

A Non-factualist defense of the Reflection principle

Stephanie Beardman

Received: 18 October 2010 / Accepted: 16 April 2012 / Published online: 19 May 2012
© Springer Science+Business Media B.V. 2012

Abstract Are there plausible synchronic constraints on how a subject thinks of herself extended over time? At first glance, Bas van Fraassen’s principle of Reflection seems to prescribe the sort of epistemic authority one’s future self should be taken by one to have over one’s current epistemic states. (The gist of this principle is that I should now believe what I’m convinced I will believe tomorrow.) There has been a general consensus that, as a principle concerning epistemic authority, Reflection does not apply to epistemically non-ideal agents. I agree with this, but argue here that it misses the point of Reflection. Rather than an epistemic principle concerning reasons for belief, Reflection concerns the semantics of belief avowal. I present a non-factual interpretation of Reflection, argue that the principle provides a constraint on the ways in which one can reflectively endorse one’s future epistemic self, and say something about the logic governing such an interpretation.

Keywords Reflection principle · Nonfactualism · van Fraassen · Rational requirement

0 Introduction

Imagine yourself, exactly four years from now. You are alive and well and working on a computer, a glass of water sitting on the desk close by. You are writing a philosophical treatise in defense of materialism, which you believe to be a true metaphysical thesis. Now, suppose that you currently believe that the above description will be true of you in four years. So, you currently believe that you will believe that materialism is true. Suppose also that this future belief conflicts with your current belief that materialism

S. Beardman (✉)
Department of Philosophy, Barnard College, Columbia University, 3009 Broadway, Milbank 326,
New York, NY 10027, USA
e-mail: sbeardman@gmail.com

is false. Never mind, if you can, on what grounds you now believe that your future belief will be as described. Here is the question for you: given your current belief about materialism, and your belief about your future belief in materialism, what should you now believe about the truth of materialism? What is the right relationship between your beliefs about your future beliefs and the rest of your current beliefs?

Bas van Fraassen has endorsed a principle he calls Reflection, which essentially says that if one is rational, one's current opinions match what one now thinks one's future opinions will be. The main criticism of this principle is that it is too optimistic—there are many cases in which it seems our future self should not, by our present lights, be believed. The literature contains a vast assortment of purported counterexamples to Reflection, ranging over cases in which one knows one's future judgments will be impaired by excessive drinking, memory loss, and so on. It seems obvious that, if I now see that in the future I will form beliefs irrationally or because of impaired cognition, I am now under no obligation to endorse those future beliefs.

These criticisms miss the point of Reflection. They arise from misinterpretations of Reflection that are both easy to make and pose an impediment to understanding what is at stake in rejecting Reflection. What is at stake is a view of the normative commitments constitutive of certain first-person belief ascriptions. Most of the misinterpretations of Reflection are the result of assessing the principle in light of the standard view of belief, which, roughly, maintains that beliefs are the sort of thing that can be justified or not by reasons (and that one has reasons for belief if and only if one has evidence for the truth of that belief). For ease of exposition, the discussion will deal for the most part with a non-probabilistic form of Reflection, which can be formulated as follows:

(Ref) Given that I now believe I will believe some proposition P at some future time, I should now believe P.¹

So, for example, if I now believe that (p) tomorrow I will have the belief that philosophical materialism is true, then I should believe now that (q) materialism is true. Because conditionals can be tricky when used intuitively, it will be helpful to keep in mind that the general form of (Ref) is Bel (q|p).

Let's start with a few descriptions of what (Ref) is supposedly about:

“Reflection can be very naturally described as mandating a kind of belief in the epistemic value of one's own future beliefs” (Christensen 1996, p. 473). He who obeys the Reflection principle is required, “in a sense, to endorse in advance his own future beliefs” (p. 459).

¹ Such a simple formulation of Reflection is used only for ease of exposition. With some further complications, the substantive comments in this paper apply to the two versions of Reflection offered by van Fraassen, which allow for both partial belief and vague degrees of belief, as follows:

GENERAL REFLECTION: My current opinion about event E must lie in the range spanned by the possible opinions I may come to have about E at later time t, as far as my present opinion is concerned (1995, p. 16).

SPECIAL REFLECTION: $P(A|pt(A) = x) = x$ That is, my current degree of belief in A, on the supposition that I will believe A at time t to degree x, should be x (1995, p. 19).

I explain and discuss Special Reflection in Sect. 3.

As an example of what Reflection requires, Maher (1992) writes, “Thus you should now be sure that you would not be in error if in the future you become sure that you can drive safely after having 10 drinks” (pp. 120–121).

Reflection “mandates diachronic coherence” (Foley 1994, p. 750), meaning by this “coherence between current opinion and future opinion” (p. 749).

The first two quotes, Christensen’s and Maher’s, suggest that, according to (Ref), one must now regard one’s future beliefs as likely to be true. The third, Foley’s, suggests that the point of (Ref) is to preserve coherence among one’s belief sets over time. A principle which has these consequences lacks *prima facie* plausibility.² Or worse: these implications are severe and epistemically distasteful enough to provide the material for an argument against such a principle by *reductio ad absurdum*. We do not think that rationality requires that we take the mere fact that a belief will be ours in the future as a reason to endorse it now, nor that it requires that we not question the reliability of our future beliefs, nor even possibly that coherence among beliefs over time is desirable. A principle which requires such things from us is bound to encounter serious resistance.

And indeed, the principle of Reflection has not wanted for critics. The quotations given above represent the voices of just a few. Section 1 provides an overview of how Reflection is almost universally understood, and discusses diachronic coherence and the extent to which one’s believing something in the future provides reason to believe it now. Section 2 evaluates the adequacy of van Fraassen’s defense of Reflection in the context of the standard interpretation. Sections 3 and 4 propose an alternative way of interpreting Reflection, one that I believe is very close to what van Fraassen intends, and one that is disengaged from the epistemic project of justification as traditionally understood in terms of reasons. There is an altogether different avenue of appeal for Reflection, having nothing to do with either diachronic coherence or the assumption that future belief plays any role in justifying current belief.

1 Some background: diachronic coherence and the inadequacies of Reflection

Much of the early literature on Reflection assumes that the value of diachronic coherence (in what I call the structural sense) supplies the motivation for accepting (Ref).³ *Structural* diachronic coherence involves an atemporal comparison of temporally-indexed (such as present and future) beliefs, to see how well they ‘match’ or agree with one another.⁴ If one’s beliefs about one’s future beliefs are true, (Ref) does, in fact, ensure such coherence between current belief and future belief. One might of course be mistaken about what in fact one will believe. But if one is committed to

² Note that van Fraassen himself, especially in his (1984) introduces the principle in a way that (misleadingly) encourages such readings (cf. pp. 243–244).

³ This was in part due to the fact that van Fraassen first presented his principle within the context of a diachronic Dutch book argument, and what seems to drive such arguments is the value of such diachronic coherence as an epistemic goal.

⁴ The greatest degree of coherence in this structural diachronic sense would involve identity between the two beliefs in question. But identical beliefs may not be the only permissible form of structural coherence; another sense in which present belief could structurally cohere with future belief is if it lies within a certain range of possible future beliefs.

maintaining such coherence as much as possible, the best one can do is to work with one's current beliefs about one's present and future beliefs.

Some writers (cf. Bovens 1995 and Maher 1992, among others) seem to think that (Ref) dictates that one should modify one's present belief to bring it in line with one's belief about the future belief. But there is nothing in (Ref) to justify such an asymmetry. Once I find myself in violation of (Ref), there are (at least) two different courses I can take to conform to (Ref): I can now believe P, or I can abandon my initial belief that I will believe P. There is, however, an interesting reason why one might interpret the conditional in (Ref) in such a way that future belief can be said to trump current belief, a reason which *would* seem to justify an asymmetrical reading of (Ref). The reason is provided by assuming Bayesianism. An agent who updates her beliefs (or rather, degrees of belief) solely by Conditionalizing on the evidence—given several idealized conditions⁵—will automatically satisfy Reflection.⁶ Within such a Bayesian context, it is natural to think that Reflection is satisfied by having future belief trump present belief. Intuitively speaking: the subjective probability assigned by a perfect Bayesian in the future would be more justified than her present probability, because more evidence would have been accumulated. Thus, if one could assume that one is a perfect Bayesian, and that one's belief system would evolve over time only according to Bayesian principles, then one would have sufficient reason to adopt the future belief rather than keep the present belief.

In any case, van Fraassen does not want his principle to commit us to Bayesianism, nor does he adopt Bayesianism (cf. his 1993). But if one rejects the Bayesian model of belief revision as an attainable ideal—if, indeed, one disregards evidence which is relevant to one's conditional beliefs—then there may be reason to think that beliefs developed later will not be as reliably generated as those formed sooner: not just in cases involving such things as addiction and delusion, but also due to a more general recognition of the effects of growing older, say, and of the psychological effects of certain experiences and situations (e.g., tendencies to rely more on habit and pay less attention to details, or to change the way one evaluates the same body of evidence based on one's emotional attachments). These sorts of cases can be developed to produce counterexamples to (Ref). With certain powerful idealizations (such as perfect Bayesianism), (Ref) is rendered redundant. Without such idealization, (Ref) is subject to numerous counterexamples.

⁵ Such as that the Bayesian's perceptual mechanisms are working well, that her beliefs are conglomerable, that there is no evil demon, that one is certain that one will always conditionalize, that one knows one's own mind, and so forth. (Roughly speaking, your beliefs are non-conglomerable when some proposition A is believed with probability less than some number x, and yet, conditional on each member of an exhaustive partition of propositions, your belief in A is greater than x. Arzenius et al. (2004) argue that Conditionalization will not satisfy Reflection when paired with non-conglomerable probabilities.)

⁶ There is a purported proof of this in van Fraassen (1995, p. 17). Weisberg (2007) shows that that proof does not work, but that an agent who conforms to Conditionalization—as long as she is certain that she will always conditionalize and always knows her own conditional beliefs and is certain that she will always have such self-knowledge—such an agent will automatically satisfy Reflection. Matthias Hild also argues against van Fraassen on this point. In his (1998a,b), Hild gives an illuminating treatment of the logical connections between Reflection and Conditionalization, stressing that Reflection does not require Conditionalization and arguing that Reflection applies all and only to (highly) idealized rational agents.

These counterexamples show that diachronic coherence, which (Ref) seems one way of trying to achieve, is not a desideratum in itself. We would reject a principle that would guarantee diachronic coherence by permitting no change at all in one's beliefs. We want to learn things, change our mind, see things differently—and it is often rational to do so. As Christensen says, “we do not think that the beliefs an agent holds at different times should cohere in the way an agent's simultaneous beliefs should” (1991, p. 241). Nevertheless, even though diachronic coherence is not itself an ideal of rationality, certain cases of diachronic incoherence might signal that things are not as they ought to be, epistemically speaking (cf. Armendt 1992). Let's look closer at this. Even in such cases, (Ref) will be the wrong remedy, at least if one is looking for reasons to justify one's beliefs.

Consider the following scenario, adapted from one of van Fraassen's arguments for (Ref). A weather forecaster currently believes, and announces on this evening's weather report, that it will not rain tomorrow afternoon. He follows a policy of updating his beliefs for the morning forecast based on data that becomes available during the night. He knows now, however, that whatever data comes in during the night, following his policies will produce the belief that it will rain that afternoon. So in the morning forecast, the forecaster will announce rain for that afternoon. Surely, something is wrong with this scenario. Such a forecaster cannot be entirely rational.

One reason to think that something is wrong is that there is no way for both today's and tomorrow's forecasts about tomorrow afternoon to be well calibrated with the objective probability of rain for that day. Moreover, the forecaster in this case *knows* that this is the case. He is knowingly committed to following procedures for updating beliefs which guarantee that his multiple forecasts about the same event (rain tomorrow afternoon) necessarily cannot each be equally well calibrated. As van Fraassen says, “It would seem irrational to organize your degrees of belief in such a way as to ruin, a priori, the possibility of perfect calibration” (1984, p. 245). Knowingly doing so would seem to impugn this person's standing as a good weather forecaster: “The forecaster's integrity, his role as a professional, is incompatible with equanimity about what has come to light” (1995, p. 16).

Van Fraassen claims that the diachronic incoherence present in this case is a symptom of a deeper defect, that the forecaster in question cannot regard himself as following a rational policy for belief revision. Something is defective in the forecaster's policy for updating his beliefs given the data that comes in during the night and the initial belief. Richard Foley (1994) adds an alternative diagnosis, according to which the forecaster's evening prediction is irrational because it doesn't take into account the evidence provided by all the possible data that will come in during the night. Like van Fraassen's explanation, Foley's “has nothing to do with the desirability of coherence between current and future opinion” (p. 761). Both explanations point to defects that lead to the formation of an irrational belief—the evening's prediction, in Foley's suggested case, and the next morning's, in van Fraassen's. Depending on how the particulars of the situation are spelled out, either of these diagnoses (or a combination of the two) could be correct.

If this is right, then van Fraassen's advice to the forecaster is startling: Given that you now believe you will believe tomorrow morning that it will rain that afternoon, you should now believe that it will rain tomorrow afternoon. In proposing (Ref) as a

solution, van Fraassen does not fix the problem, but only addresses the symptom. It solves neither the problem identified by Foley nor that posed by van Fraassen himself. One sort of defect could be remedied by adopting a different policy for updating beliefs, the other, by taking into account all the relevant available evidence.

The moral that emerges from the discussion so far is this. Becoming diachronically coherent in the structural sense won't solve the forecaster's problem. What matters instead of such coherence—and what the presence of structural diachronic incoherence *might* indicate, as it does here—is (a) whether one is currently rational (e.g., whether one has properly attended to all the available relevant evidence in deciding what now to believe), and (b) whether one now can endorse one's future changes in belief as being the result of a rational process of belief revision. The latter points to the importance of a second sense of coherence. According to the *dynamic* sense of diachronic coherence, two beliefs at different times are diachronically coherent if and only if there is a 'rational' (in a sense which needs to be defined) evolution from one belief to the other. What we care about, epistemologically speaking, is dynamic diachronic coherence, rational belief formation over time.

Given this analysis, Foley and others have offered promising proposals to replace Reflection.⁷ I take his proposal—or something very like it⁸—to yield the right answer to the question that everyone thought (Ref) was trying to answer: What role should one's future beliefs play in one's current deliberations about what to believe? This is so because it respects what we intuitively think is the right connection between evidence and belief—namely, that one has reason to believe P if one has evidence that P is true. We needn't go into the details of these proposals. What's relevant here is that they have helped to solidify a diagnosis of what is wrong with (Ref). For epistemically non-ideal agents, the principle seems to enjoin us to go beyond evidentialism, beyond the

⁷ The following principle, proposed by Foley (1994, p. 762), captures what is initially appealing about Reflection without implying, as Reflection does, that structural diachronic coherence is intrinsically good.

The Rational–Rational Generalization: *Rational* current opinion coheres with future opinion as seen by *rational* current opinion.

This is supposed to be a rough generalization. There will be many cases—such as those which are offered as counterexamples to (Ref)—of which this generalization is not true. Foley provides the following rule for determining when the generalization obtains. (He further argues for three theses to flesh out this rule, cf. his 1994, p. 762; and 2001, p. 131ff.)

Rule: The rational-rational generalization obtains when my reasons for thinking that I will believe P also provide me with evidence for P itself; otherwise, it doesn't.

If my reasons for thinking that I will believe P also provide me (now) with evidence for P itself, then rational current opinion will as a matter of fact cohere with future opinion as seen by rational current opinion. If, on the other hand, my reasons for thinking that I will believe P do not also provide me with evidence for P itself, then this is just the sort of case in which rational current opinion should *not* cohere with what one currently believes future opinion to be.

⁸ Another proposal that is in keeping with the direction laid out by Foley is Luc Bovens' 'diachronic constraint on rational belief': "I propose that in determining whether we have sufficient reason to believe that P, we should be impartial with respect to the time at which the evidence was, is, or will be available" (1995, p. 755). In other words, whatever evidence for P one now believes one has (timelessly) sufficient reason to believe, gives one now sufficient reason to believe P. The epistemic criteria—evidence and sufficient reason for belief—are timeless in the sense that they're seen as always relevant to the goal of now believing truths and now not believing falsehoods.

constraint of relying upon evidence alone in deciding what to believe. (Ref) would seem to be advocating non-evidential reasons for belief, and to have us believing something in advance of the evidence for it. The mere fact that you will believe P does not in itself count as *evidence* for the truth of P, and thus, it is not in itself a reason to believe P now.

It is natural to think of (Ref) in terms of future epistemic authority. This is the perspective adopted throughout nearly all the relevant literature. From this perspective, (Ref) provides an (incorrect) answer to the question of what epistemic relevance one's future beliefs should have to one's current beliefs, from one's current vantage point. But this is not the right way to construe what the principle is about. There is something importantly right about Reflection, and this something cannot be discerned by adopting the usual assumptions of its critics.

As we will see, acceptance of (Ref) has nothing to do with the epistemic authority of the future self. (Ref) can only be a constraint on the present self, and there is no reason the present self must presume that in general her future course will be one of epistemically superior development. It is not best understood in terms of treating one's future self as an 'expert'. This is likely to be missed if one is thinking in terms of 'reasons for belief', and even more so if one is thinking in terms of probability functions that are formulated, as they usually are, from the point of view of an external observer.

2 Personal identity and integrity

Before turning to the new interpretation on offer, let's look at what's wrong with the standard reading of van Fraassen's own defence of Reflection in response to the widespread criticism referred to above. In what he calls the 'death and disability defense', van Fraassen claims that the principle is to apply to all and only states that the agent takes to be *genuinely* hers. "In Reflection," he writes, "I refer to the range of genuine values and opinions which are genuinely mine, at the relevant future time, as far as my present opinion of that future allows – not to anything I do not classify as such" (1995, p. 24). So the question, in his words, is "what I am willing to classify as future opinions of mine" (p. 23). According to this defense, the purported counterexamples involving future drunkenness, brainwashing, forgetting, and the like, do not apply because in such cases the future self is not properly classified as one's genuine self.

The death and disability defense cannot be correct if taken literally. (I do not mean to suggest that van Fraassen himself intends it to be taken literally.) There are three main reasons to reject it. *First*, my identity is not imperiled—nor do I take it to be—when I have good reason not to satisfy (Ref). It's implausible to think that it is a metaphysical condition of diachronic identity that one satisfy Reflection. For example, it's not plausible that one must think one will not be the same person when drunk, when uncertain about whether one has been drugged, when sailing past the Sirens, and so forth. An agent needn't think that a dissolution or break in personal identity is imminent when (Ref) does not apply. Correspondingly, having a self that one takes to be unified over time does not require that one believe now that one will change one's beliefs in ways that are rational, or that would justify trusting one's future self.

Second, to be able to ascertain which of my future beliefs are ‘genuinely mine’, so that one could determine when the principle of Reflection is applicable, it seems one would need some hint of a principled way of determining which future states are ‘genuinely mine’, which a reading of van Fraassen’s work does not readily provide. For example, we might propose that which future opinions are mine is plausibly determined by what I deem pathological or not. Or, similarly, we might try adapting Jeffrey’s (1988) suggestion, and say that x is a future opinion that’s genuinely mine whenever I can safely assume that the opinion in question will be formed by a series of reasonable steps. Finally, we might pursue van Fraassen’s own suggestion, that states are genuinely mine only if I think they are formed “under my control” (1995, p. 26).⁹ (He rejects the first two proposals because he wants Reflection to be *violated* when I think I will become unreasonable or pathological, and not merely rendered inapplicable.) Without dwelling on the involved and interesting details of how these proposals might be fleshed out, it will be sufficient for our purposes here to point out that none of them will be able to provide an adequate defense against an array of purported counterexamples. Each of these proposals attempts, with varying success, to address the fact that future opinions are often rationally un-endorsable. But there is a class of counterexamples that do not have that feature. Rather, the future state in these cases *can* be rationally endorsed (for the future self) from the present perspective. These cases involve neither epistemic unreliability nor the absence of being in the proper sort of control—and yet they appear to be cases in which Reflection seems rationally violated.¹⁰

There seem to be many familiar and apparently widespread examples of this phenomenon. I take the clearest of these to involve evaluative shifts, but there may be other types as well. These involve reasonable, predictable shifts in how the same body of evidence is weighed, given a change in epistemic context generated by new circumstances and attachments. For example, one might foresee that once one becomes a parent, some of one’s current beliefs will undergo a radical shift simply because of a change in the perspective from which one evaluates a given body of evidence (say,

⁹ More recently, Simon Evnine has endorsed a version of Jeffrey’s criterion for when the future self should be treated as an expert (arguing for a non-circular way of identifying what is to count as ‘reasonable’), combined with an expansion of van Fraassen’s control criterion of identity. According to Evnine’s notion, my being in control of my behavior depends on not needing to “externalize my current intentions and knowledge” in someone or something else in order to ensure that that behavior takes place (2007, p. 107). One is in control if one’s present intentions are sufficient to bring about one’s future self’s behavior. Evnine recognizes that this is too strong: surely there is no breakdown in personal identity just because I happen to need to write something down in order to get myself to do it later. He suggests that a breakdown would occur only when the extent to which I cannot control future actions merely by intending them becomes “too great”—he speaks in terms of taking “serious measures”, like handing over power of attorney to another.

¹⁰ Take, for example, the strange violations uncovered by Arzenius et al. (2004), in which one conditionalizes on the evidence and yet rationally violates Reflection because one’s probability functions are not conglomerable (see footnote 5 above). The argument for this, and for the rationality of non-conglomerable probabilities itself, is based on several decision-theoretic assumptions that we need not get into here. The import of that paper for present purposes is to raise the possibility that there may be cases involving perfectly rational belief shifts which lead to apparent Reflection violations. Note also that even *uncertainty* (about whether some sort of cognitive impairment will occur) can justify violations of Reflection (cf. Elga 2000). On my view, this is because such uncertainty can justify the withdrawal of identification with the future self.

about whether the suburbs are a desirable place to live). Or, take the oft-cited example of one who recognizes that they are likely to become more politically conservative as a result of gaining wealth or growing older. Likewise, we often recognize that undergoing certain educational and professional training and the like will have an enormous impact, often in predictable ways, on the sort of person we will become—not merely for the obvious reason that we will gain new information, but because the experience can cause us to change current beliefs even without introducing new evidence or raising our epistemic standards. In these cases, it can be rational to endorse the change in the future belief in the context of those future circumstances, without endorsing that belief now while one is not in those circumstances. The divergence between present and future opinion does not indicate any sort of unreliability (nor need it indicate an improved epistemic perspective), from one's present perspective. Though I don't (and shouldn't) believe it now, I can nevertheless endorse that my future self believe it in different circumstances. What to make of such cases, and whether they have been properly characterized here, are controversial matters. I mention them because they seem to count as cases in which the future self is 'genuinely me' (by my present lights), and yet it is rational to violate Reflection. The future self here is simply me in a new circumstance—after having children, getting wealthy, joining the military, or whatnot.¹¹ So there appears to be a class of counterexample that survives the death and disability defense.

Third, putting things in terms of 'death and disability' is a mistake, and crucially misses the central import of Reflection. This is not only because doing so invites questions about the criterion of personal identity, but because it invites the misreading of (Ref) that has produced such a large literature. It becomes easy to think that whatever criterion is in play will be identifiable from a third person perspective, and this is to head in exactly the opposite direction needed for understanding (Ref). These strategies of restricting Reflection, based on the manner and extent to which future selves can be considered experts or be considered the same person, *might* offer promising answers to the questions concerning what we have reason to believe, and what epistemic weight a future self should have. But those are not the questions that Reflection itself is meant to address, and answering them does not provide any basis for understanding and defending what is important about Reflection. Rather than focusing on the method by which the future belief is acquired, and whether or not it is reliable, it seems to me better to invoke something like 'what I am now willing to commit to as an expression of myself'. And this will not be determined by questions of epistemic reliability, nor by the likelihood that I will remain relatively free and in control of my mental states. Instead, satisfying Reflection is itself integral to one's first-personal, present-based, expression of one's agency. It is not possible coherently to commit to a violation of (Ref). Under this interpretation, talk of future selves is inappropriate. My belief that 'P and I will believe P', when understood as I argue in the next section that it should be, conceptually ties together, as it were, the notion of separate selves. One is just one self, unified in time, *from the perspective of the present believer*. Which future states

¹¹ The sort of case described here is ubiquitous, and deserves further study. The examples in this paragraph were generously presented to me (albeit in a different context) by an anonymous referee.

are ‘genuinely mine’ is not a question of identity over time, but a matter of what I am willing now to stand by as an expression of my current self.

The rest of the paper will urge a shift in perspective from the traditional epistemic concerns of this section. My own proposal, described in the next section, eschews the project of providing an account of personal identity that might be identified from the third person perspective. Nor does it rely upon a notion of endorsing the relevant future state. Rather, Reflection has to do with a pragmatic synchronic coherence, the maintenance of which entails that the question of the future self’s epistemic authority is a moot one. I will argue that (Ref) is a constraint on what it takes to avow a doxastic conditional commitment to a proposition. The roots of my defense are evident in van Fraassen’s writings on the matter, and much of what follows can be read as a sympathetic extension of what he may have had in mind in proposing (Ref) as a synchronic constraint on rational thought. What most of his critics haven’t noticed, or haven’t taken seriously enough, is that van Fraassen has changed the question. Reflection is a terrible answer to the question of how much epistemic authority to grant the future self (by one’s present lights) in part because it’s not intended to answer that question; it doesn’t address the question of what provides adequate reason for belief, and it doesn’t claim that my believing I will believe P provides me with a reason now to believe P. The main criticism of (Ref), that it insists upon a non-evidential reason for belief, is missing the point.

3 The non-factualist interpretation of Reflection

The defense is inextricably tied to the present tense, first person point of view. This is a point the importance of which the literature has generally ignored. Hild (1998a) explicitly acknowledges and defends this widespread practice (of ignoring the first person perspective) when he argues that the first person point of view can be “harmonized with the observer’s point of view. We could easily,” he says,

...shift back into talk from the point of view of an external observer and interpret the reasoner’s auto-epistemic views about her own epistemic states as the observer’s views about the reasoner’s epistemic states. We could also think of the reasoner’s auto-epistemic views as an echo of the observer’s views from whom she could have borrowed her views about her epistemic situation (p. 237).

While the equivalence between first and third person points of view is true for Hild’s hyper-idealized agent, it is *not* true if one wants to defend (Ref) for non-ideal agents, as I do.¹² In the case of non-idealized agents, the only defense possible of (Ref) will be from the first person, present tense, engaged and expressive point of view of the agent. It cannot be defended from the perspective of the disengaged observer.

¹² It’s possible that Hild would agree with me about this, though he does not address the point. Rather, he presents a defense of Reflection that “consciously confines itself to ideally rational agents” (1998a, p. 230).

An attempt to formalize this engaged element is evident through the use of capitals and lower-case symbols in van Fraassen's Special Reflection.

Special Reflection $P(A|pt(A) = x) = x$

We can read this as follows: 'I now take it to be likely to degree x that A , on the supposition that it will seem likely to me to degree x that A at time t .'¹³ According to van Fraassen,

the function of the sentence $P(\dots) = x$ as a whole is to express the epistemic attitude of self-attribution, while the contained $pt(\dots) = x$ states the fact that (at t) I have this attitude. In English, of course, both roles are played by the same words (1995, pp. 8–9).

We must understand the 'I now take it to be likely...' in my gloss of Special Reflection to govern the whole sentence, and to express something about myself. The notion here is that of expressing—or better, avowing—a subjective conditional probability. The function of the sentence as a whole is to express an epistemic attitude, while the embedded sentence $pt(A) = x$ states a fact about the speaker at time t .

This reading applies to our qualitative interpretation of a statement which satisfies (Ref) as well. Thus 'I will believe that A ' is a statement of autobiographical fact; not necessarily so for the present-tense statement 'I believe that A '. One who avows the latter is doing something, expressing what van Fraassen identifies as a kind of commitment. The relevance of this distinction to the present discussion is that the avowal/expression of belief involves normative commitments that are different from those of belief reports or descriptions. In particular, the commitments of belief avowals render certain (factually plausible) combinations of expressions and reports pragmatically incoherent.

¹³ I have taken the liberty of modifying van Fraassen's own interpretation so that it is explicitly in the first person. Van Fraassen construes sentences of the form $P(\dots) = x$ as 'It seems likely to me to degree x that I am such that...' (1995, p. 8). He allows that "[t]he phrase 'that I am such that A ' may indeed generally be replaced by 'that A ' in the reading without loss" (1995, p. 8). Accepting this shorthand, Special Reflection can be read as follows: 'On the supposition that it will seem likely to me to degree x that A at time t , it now seems likely to me to degree x that A .' Van Fraassen makes clear, however, that "the important point" is that "the sentence as a whole is in the first person singular present tense indicative, while of course no such restriction applies to the contained sentence A " (1995, pp. 8–9).

For van Fraassen, personal probabilities are not propositional attitudes, but rather, are modes of attributing certain properties to oneself. That is, it is not correct to attribute to one who says 'I believe that A ' a relation to a proposition [that A]. Rather, the correct interpretation is that the speaker is performing an act of self-attribution. In this way, the belief statement just referred to is understood to be saying 'I believe myself to be such that A '. The view endorsed here by van Fraassen was developed by David Lewis (1979). As an aside, it is significant that the view grew out of problems that arise in trying to ascribe propositional attitudes to self-referential statements of belief (such as, 'I believe that I am F '). Analogously, it is the self-referential aspect of the sort of statements to which (Ref) applies (i.e., those that make reference to what one will believe in the future) that necessitates the sort of defense marshaled in this paper. However, van Fraassen needn't have adopted Lewis's interpretation of subjective probability as a mode of self-attribution in order to secure the elements needed here. The better, and more salient, distinction is that between reporting and expressing one's intentional states.

Van Fraassen amplifies what is involved in what he calls expressions of belief as follows:

I conclude that my integrity, qua judging agent, requires that, if I am presently asked to express my opinion about whether A will come true, on the supposition that I will think it likely tomorrow morning, I must stand by my own cognitive *engagement* as much as I must stand by my own expressions of commitment of any sort. ...But I can no more say that I regard A as unlikely on the supposition that tomorrow morning I shall express my high expectation of A, than I can today make the same statement on the supposition that tomorrow morning I shall promise to bring it about that A. To do so would mean that I am now less than fully committed (a) to giving due regard to the felicity conditions for this act, or (b) to standing by the commitments I shall overtly enter (1984, p. 255).

Like making a promise or expressing an intention, avowing a belief does not consist in making a factual report. Rather, like these other activities, avowing a belief invites others “to rely on my integrity and to feel assured that they now have knowledge of a major consideration in all my subsequent deliberation and courses of action” (1984, p. 254). Let’s call the mode of construing first person belief attributions which recognizes an element of commitment ‘non-factualist’. As noted above, one can also adopt a third person attitude towards one’s own beliefs; call this attitude ‘factualist’. According to the distinction then, the *factualist* makes autobiographical statements about her beliefs. These statements have much the same status as third person statements have. The *non-factualist*, on the other hand, is not making a statement of autobiographical fact when she states what she believes—rather, she is avowing a commitment.¹⁴

Van Fraassen claims that satisfying Reflection involves “the commitment to proceed in what I now classify as a rational manner, to stand behind the ways in which I shall revise my values and opinions” (1995, p. 26). It involves not only the commitment to change one’s opinion for good reasons (to reach the future with a series of rational steps), but to take responsibility now for one’s future opinions. This is essentially a commitment to stand by one’s commitments (1984, p. 256). On my reading, the relevant commitment is not to the future belief—it is to a present belief, conditional on future belief. The semantic logic of the sincere expression of belief conditional upon a statement of future belief is different than the logic of disengaged, third-personal reporting of such conditional belief. All doxastic avowal is agential, and involves doing something that incurs obligations. This is why one cannot express (in the sense of avow) future commitment. In resolving to “stand behind the ways in which I shall revise my opinions”, I am resolving to take myself to have reached the future with a

¹⁴ Alvin Plantinga is one of the few philosophers to embrace this point. He likens the general idea to what is involved in a confession of faith: “to confess your faith is not (or not merely) to make an autobiographical statement about the condition of your psyche” (1993, p. 152). For example, “if I use [the Apostles’ Creed] to express what I believe, I am not merely reporting the result of self-examination, as I might be if I told you that every now and then I am subject to doubts about one or another element of the creed. I am doing something quite different: something that involves making or renewing a commitment” (p. 152). He thinks there is an element of commitment, required by one’s integrity as a believing agent, involved in all belief avowals. “Seriously using a creed...,” he explains, “is only a special case of a more general phenomenon: stating or expressing one’s considered opinion” (p. 152).

series of rational steps. Furthermore, I can resolve to do so only for as long as it *is* possible to see myself in this way, and no longer.

On my reading, the commitment involved in Reflection is a commitment to see oneself in a certain way. It is a commitment to see oneself now as changing one's opinions rationally, absent any evidence to the contrary; this is the default position in how one sees oneself, going forward. An epistemic agent with integrity would have to make assumptions that conform to Reflection. This identification with one's future self is what enables, and thereby requires, that I am now, as van Fraassen puts it in the quotation above, "standing by the [doxastic] commitments I shall overtly enter," while giving "due regard to the felicity conditions for this act." But this last clause implicitly acknowledges that there are conditions that would obviate the commitment; Reflection is a present-tense conditional avowal, an avowal of conditional commitment.

To see this point, return to the example of the weather forecaster. This is the starting point of the argument for Reflection: something is incoherent—unintelligible—in cases like that of the forecaster. As we've seen, his problem is not a matter of maintaining diachronic coherence, that is, of getting his present and future beliefs to (structurally) 'cohere'. Rather, his problem is that he has reason to believe that he will have a certain belief that conflicts with another belief he now has. He is faced *now* with *believing* that his beliefs are not diachronically coherent. As explained earlier, this puts him in a predicament. He must admit that either (a) his current belief about the chance of rain tomorrow is irrational because he has not taken into account all available and relevant evidence for the chance of rain, or (b) his own policy for updating his belief based on the evidence that comes in during the night is itself irrational, in which case he is now not committed to a rational process of belief revision. Either way, the violation of (Ref) indicates that irrationality is present in the forecaster. What makes the situation incoherent is that the agent is admitting that he has *now* abdicated his responsibility as a believer. The dilemma he is faced with when violating (Ref) undermines his integrity as an epistemic agent.

Let's explore this epistemic predicament further. What happens if the description of the situation is altered to include an explanation for why the forecaster's morning statement (about the probability of rain tomorrow) will differ from his current forecast? As before, he knows that no matter what data comes in during the night, he will update his beliefs so that his morning forecast predicts rain, while he currently predicts no rain for that day. Let's add to this a story about why there is such a discrepancy in the forecasts—maybe that he will get exceedingly drunk tonight, and knows that whenever he has a hangover the next morning he predicts rain. Or choose whatever sort of story you prefer to add here You might now reasonably object that my proposed interpretation of Reflection does not present the right take on epistemic integrity. You might ask, with some incredulity: have I really abdicated my epistemic responsibility if I foresee that I will be brainwashed, drugged, or seduced into having beliefs that I disavow now? On the contrary, you would say, if I knew I was facing such a future, my integrity as an epistemic agent would seem to require me to violate Reflection. So it looks like the standard counterexamples to Reflection apply here too, to the non-factual version. A more general way to put the worry is this. Let's say it is true that the commitment to take oneself to be a reliable epistemic agent applies in some relevant sense to *present* beliefs—something like that may be necessary in order to count as

having beliefs at all. But why must this sort of commitment apply to my beliefs about my *future* beliefs? My future beliefs are not (and cannot be) avowed now. Why must I now be confident of *them*?

To answer this challenge, we have to investigate what difference non-factualism makes. As a preliminary, and to help forestall misunderstanding, let's reiterate an obvious point. Any principle that yielded the result that one couldn't possibly rationally believe that one would become brainwashed, become an unreliable epistemic agent (by one's present lights), forget things, evaluate evidence differently, or adopt different standards of rationality than one accepts now—such a principle would have to be ruled out at the outset. One can know that one's opinions will change, even that they are likely to change, in ways that result in conflicting present and future beliefs. There is nothing irrational or incoherent about knowing this. Indeed, this is exactly why Reflection doesn't work as a way of identifying one's reasons for belief, as argued above in Sect. 1. One can coherently assess the likelihood that one's subjective probabilities match up with the objective probabilities. In doing so, we disengage from, or stop identifying with, our future beliefs. (Ref) does not preclude this because (Ref) is not meant to govern factualist belief ascriptions. As van Fraassen himself admits, Reflection "is in fact indefensible if we regard the epistemic judgment – whether formulated in probabilistic or more qualitative terms – as a statement of autobiographical fact" (1984, p. 256). It is possible to step out of the mode of commitment and self-expression, to rationally assess one's likelihood of being wrong about oneself. But when doing this, one is not avowing a commitment but viewing one's opinions descriptively.

This means that, like all conditional commitments, the commitment expressed by Reflection is defeasible. It is defeasible, for instance, by evidence that one will not follow a rational policy of belief revision. So one *can* say 'I believe P and I will believe not-P'. But one *cannot* (coherently) do this without some intelligible reason for the change in opinion (and though it must be a reason of a certain type, as explained in the next paragraph, it may be enough to believe that such a reason obtains, without knowing exactly what it is). The case of the weather forecaster bears this out. Absent some story (like drunkenness, brainwashing, or the like), the opinion of the forecaster is simply incoherent. Something has to be wrong with it. Consider this. If you're in the position of the weather forecaster, it is not an option to point out sheepishly, as a way of trying to argue that (Ref) does not apply in your case, "Well, *if* I knew I would be brainwashed tonight, my beliefs *would* be coherent", when you don't think that that condition is likely to obtain. Absent some attending explanation (or the belief, based on the testimony of a trusted authority, that such exists), the default in belief statements is non-factual, and the relevant commitment is the one encoded by (Ref).

The sort of supplemental explanation required to make sense of an apparent violation of (Ref) is one that effectively takes the conditional component out of the scope of the first-personal avowal. Such an explanation marks the enterprise as one of evaluating evidence and reasons for belief. It pulls the prediction out of the scope of one's present avowing, and holds it up as a factual report to be examined from a position that is identical to one that a third person could undertake. In these cases, the antecedent of the conditional will be rendered detachable. Thus, rather than $\text{Bel}(A|B)$, the logical structure of the sentence would be $\text{Bel}(A)|B$. It seems significant that such third-personal explanations seem necessary for understanding any sort of statement of the

form ‘I believe P and I will believe not-P’. These explanatory reasons may be implicit and contextually-defined, and they need not always explain failures of rationality. For example, they can range from cases in which one is confident one will forget something one now believes, to cases in which one foresees one will feel differently about one’s current valuations, to cases of cognitive impairment of every sort. Seeing what is required to explain apparent violations of (Ref) implicitly acknowledges the legitimacy of (Ref). The only sincere conditional expression of a commitment that’s coherent is one in which the agent identifies with the condition in a way that makes it part of her current avowal.

There are two different frameworks, or perspectives, at play. According to one, you need to know how and why your future beliefs change in order to know how to assess the case epistemically. That is because you want to evaluate whether there is good (evidential) reason for the change, and what sort of reason it is. This is an enterprise conducted in terms of justifying reasons for belief, and it is the one that has been assumed in nearly all of the discussion about Reflection to date, as discussed in Sect. 1. But there is something else here, too, which is the subject of the present section: the non-factual reading of belief avowals. If there is no third-personal explanation available to create distance and pull a statement out of the scope of one’s avowals, and one avows simply ‘P and I will believe not-P’—then something is clearly wrong. (Ref) is violated, and we find the case unintelligible. We cannot take such a person seriously as an epistemic agent.

In sum, foreseeing a change in, or even a failure to live up to, one’s present epistemic commitments does not (necessarily) amount to abdicating one’s responsibility as an epistemic agent. Coherently foreseeing a failure in epistemic reliability by one’s present lights forces one into a non-engaged, non-agential conception of one’s self. One cannot do it from within the first person present expressive point of view, without becoming incoherent as an epistemic agent. In practice, we can usually tell the difference by whether some sort of explanatory (usually exculpatory) story is available to explain the shift from avowal to belief report. Moreover, in overriding the first person present tense conditional obligation, the acknowledgment of such an explanation changes the logic of one’s avowals—the future belief is no longer within the scope of the main operator (the non-factual avowal). Absent a special story, the default is an engaged condition of semantic coherence.

4 Reflection as a normative requirement

It is important to see that the non-factual, ‘voluntarist’ position defended in this paper is not a form of non-evidentialism. Non-evidentialism is the view that there can be reasons for believing that for which one has insufficient evidence: some facts provide inadequate evidence for the truth of a belief, and yet provide adequate reasons for having that belief. I take the evidentialist/non-evidentialist distinction to concern itself only with the relation of evidence to having reasons for belief. But when first person present tense belief ascription is construed non-factually, the question of whether it is justified by reasons does not come up. As a form of “cognitive commitment, intention, engagement” (1984, p. 256), it is not the sort of thing that could be justified

on a purely evidential basis. This type of voluntarism is not a sort of non-evidentialism in the sense that it does not provide a non-evidential reason for belief. Speaking intuitively, there is a gap between what provides adequate reason for belief and what constitutes rational belief. Van Fraassen embraces this point when he both endorses a strict, empiricist, reading of evidentialism *and* he maintains that “[i]t is not irrational to ‘go beyond the evidence,’ and belief in angels or electrons or the truth of theories in molecular biology does not *ipso facto* make one irrational” (1985, p. 248). As he puts it, “rational belief change is not restricted to mere accommodation of what one takes as evidence” (1985, p. 250). So, rather than doing epistemology in the sense of being concerned with justifying reasons for belief, we are here making a contribution to the logic, or semantic analysis, of belief avowals. Reflection is a semantic constraint on a non-factual avowal of one’s self-understanding. Questions about evidence, and about reasons for belief, are not part of the analysis of what it is that one is doing when one avows a belief.

John Broome has argued that the reason that certain requirements of rationality do not tell us what reasons we have is that these requirements—he calls them normative requirements—have the formal property of being wide-scope. *X normatively requires* you to ϕ iff you ought (to ϕ if X). The brackets are needed here to make it clear that the scope of the ought operator is wide, and to distinguish it from the narrow-scope requirement according to which ‘if X, you ought to ϕ ’. For example: “Believing the world was made in six days normatively requires you to believe it was made in less than a week” (2004, p. 29). But it is not true that if you believe the world was made in six days, then you ought to believe it was made in less than a week. This is because you could detach from this narrow-scope ought, via modus ponens, the false normative conclusion that you ought to believe the world was made in less than a week. But it’s still true that you ought (so to believe if you think the world was made in six days).

Reflection has the form of a wide-scope ought. It says: you ought (to believe P given that you now believe you will believe P at some future time). No normative conclusion can be detached from this requirement; one cannot derive from it the conclusion that you ought to believe P or that you have a reason to believe P. As a wide-scope requirement of rationality, Reflection has the logical structure of a normative requirement.

However, because of its non-factualism, Reflection is strikingly different from the collection of normative requirements discussed by Broome and his interlocutors.¹⁵ Though there isn’t the space to fully develop these now, it may be worth noting here a few features that possibly suggest the interest and significance of recognizing a special logic of non-factual requirements. First and most obviously, (Ref) is a semantic constraint on intelligibility, and not a practical norm of reasoning. As such, some will think that the principle is not a normative one at all. There is already debate over whether and to what extent ‘normative requirements’—wide-scope oughts—are in fact normative. [Most recently, in the debate between Kolodny (2005) and Broome (2007). Earlier, these issues have been discussed in the deontic logic literature, see for

¹⁵ Here is a sample of the sort of requirements typically invoked. (a) Intending an end normatively requires you to intend what you believe to be the necessary means to that end; (b) Believing you ought to ϕ normatively requires that you intend to ϕ ; (c) The fact that you believe p, and that q follows from p, normatively requires that you believe q.

example the view of conditional obligation presented by [Loewer and Belzer \(1983\)](#).] The question of whether wide-scope oughts are normative becomes even more pressing if Reflection is to be counted among such requirements. Of course, one might think semantic constraints should be excluded from debates concerning the nature of normative requirements.¹⁶ Broome, for example, is mostly interested in norms of reasoning, and Reflection has nothing to do with such norms. He may be content to confine his arguments to norms of reasoning, both theoretical and practical, and so he may be happy to find some way to restrict his logic and semantic models to requirements that govern those activities. On the other hand, recognizing Reflection as a (semantic) normative requirement might help shed light on the general notion of conditional obligations, and on such phenomena as conditional promises and conditional intentions.

Second, Reflection, as expressed by (Ref) and as made precise in Special Reflection, cannot be represented by the material conditional. Using the material conditional would overlook the defeasibility conditions for Reflection.¹⁷ Third and relatedly, Reflection offers an example of a principle in which the wide-scope of its logical form is essential to it. This may have implications for the Broome–Kolodny debate concerning whether the term ‘rationality requires’ must take wide or narrow scope. For example, in his (2007), Broome offers a proof that the property of being rational is independent of whether normative requirements have wide or narrow scope. “The proposition that you are rational is the same whichever way a conditional requirement is formulated,” he says. “Either way, you are rational at exactly the same worlds” (2007, p. 363). But this is not true if we include Reflection among the normative requirements. If constraints on non-factual avowals are admitted as normative requirements, the logical terrain seems to shift markedly. Reflection would then be an example of a requirement whose wide scope is essential to the question of whether things are both epistemically and semantically as they should be. This is because the wide scope of Reflection is secured, as it were, by the non-factual element; there is simply no corresponding narrow scope formulation. Certain non-factual requirements, in other words, are intelligible only as wide-scope requirements. This suggests that the debate between wide and narrow scope may play out quite differently if non-factual requirements were to be considered.

These three matters require further investigation. I mention them here to point to possibly fruitful areas of research that emerge once Reflection is accepted. What I

¹⁶ Because I argue that (Ref) is a semantic constraint on intelligibility and also claim that it is a ‘wide-scope ought’ in Broome’s sense, one might worry that I am running together normative, epistemological, claims with descriptive, semantic, ones. However, the distinction in this context is not as clear-cut as one might hope, since there is not a consensus on how to understand wide-scope oughts. They are not obviously normative in the sense that are narrow-scope oughts. At the very least, of course, they do not yield practical reasons.

¹⁷ Since he is mostly interested in capturing the non-detaching feature of wide-scope oughts, Broome could keep that feature and grant that the material conditional, which he uses, is not essential to his account. In a footnote, he allows that the notion of conditional obligation used in deontic logic, $O(q | p)$, could provide a formal model for the notion of a normative requirement (2000, p. 82). However, though he notes that ‘ $O(p \rightarrow q)$ ’ is not equivalent to ‘ p requires q ’, he seems committed to the claim that the material conditional is a consequence of it. And indeed, his (2007) argument concerning the equivalence of wide and narrow formulations regarding the property of rationality depends on the use of the material conditional.

hope to have established in this paper is that there is some reason for thinking that the logic of non-factual requirements may differ from that of factual requirements.

Understood as a principle governing a factualist interpretation of belief ascription for non-ideal agents, (Ref) is indefensible. Non-factualism is the best way to understand and appreciate what is involved in Reflection, as a condition of epistemic integrity and self-commitment. Does (Ref) preclude us from being pessimistic about our future epistemic selves? Can we question the truth of our beliefs—in this case, of our future beliefs? Of course we can. But not in the same breath, as it were, that we use to avow our beliefs. (Ref) is unique as a requirement of rationality in that it codifies a requirement of *seeing oneself as rational*; it's a constraint on how one thinks of oneself, given a non-factualist construal of belief avowal. A factualist interpretation allows for critical evaluation of the belief in question, inviting doubt about its truth. In contrast, a non-factualist interpretation forecloses all questions of doubt, and requires a confident endorsement of one's beliefs. We shouldn't argue that one attitude is more rational than the other; there is a different logic governing each. Say you're a self-assured anti-materialist. Should you believe that materialism is true on the hypothesis that you will think so in a year? If you think the predicted future is truly inescapable, and you cannot endorse your future belief because you have no reason now to think that it will be reliable, you must either dissociate from your future self or revert to merely reporting your current belief. What you cannot sincerely say, on pain of incoherence, is a statement that violates Reflection. You cannot avow, non-factually, the belief that: materialism is false and you will believe it to be true. Essentially, you cannot now promise to abdicate your epistemic responsibilities as you presently construe them. One might put it this way. In expressing our epistemic commitments, we must—in that very moment of self-expression—operate with the assumption that we are ideal epistemic agents, while fully intending to earn that faith. We always knew that Reflection applies to idealized epistemic agents. What's interesting here is that it applies to us, too.

Acknowledgments For many helpful discussions, I am grateful to John Collins, Frankie Egan, Dick Foley, Brian Loar, and Barry Loewer. I am also grateful for a reading group discussion of my work with Matt Evans, Liz Harman, Nishi Shah, and Sharon Street; the comments of two anonymous referees for this journal; a long ago conversation with Bas van Fraassen about Reflection; and for the many comments on nascent versions of this work given by my peers in graduate school at Rutgers, (especially Elisa Galgut, Orit Gwirceman, Justine Kingsbury, Andrew Milne, Anna Vitola, Julie Yoo, and the members of the proseminars at which I first presented my ideas on Reflection).

References

- Armendt, B. (1992). Dutch strategies for diachronic rules: When believers see the sure loss coming. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1, 217–229.
- Arntzenius, F., Elga, A., & Hawthorne, J. (2004). Bayesianism, infinite decisions, and binding. *Mind*, 113(450), 251–282.
- Bovens, L. (1995). 'P and I will believe that not-P': Diachronic constraints on rational belief. *Mind*, 104, 737–760.
- Broome, J. (2000). Normative requirements. In J. Dancy (Ed.), *Normativity* (pp. 78–99). Oxford: Blackwell.
- Broome, J. (2004). Reasons. In R. J. Wallace, P. Pettit, S. Scheffler, & M. Smith (Eds.), *Reason and value* (pp. 28–55). Oxford: Clarendon Press.

- Broome, J. (2007). Wide or narrow scope?. *Mind*, 116(462), 359–370.
- Christensen, D. (1991). Clever bookies and coherent beliefs. *Philosophical Review*, 100(2), 229–246.
- Christensen, D. (1996). Dutch-book arguments deprivatized: Epistemic consistency for partial believers. *The Journal of Philosophy*, 93(9), 450–479.
- Elga, A. (2000). Self-locating belief and the sleeping beauty problem. *Analysis*, 60(2), 143–147.
- Evnine, S. J. (2007). Personhood and future belief: Two arguments for something like Reflection. *Erkenntnis*, 67, 91–110.
- Foley, R. (1994). How should future opinion affect current opinion?. *Philosophy and Phenomenological Research*, 54(4), 747–766.
- Foley, R. (2001). *Intellectual trust in oneself and others*. Cambridge: Cambridge University Press.
- Hild, M. (1998a). The coherence argument against conditionalization. *Synthese*, 115, 229–258.
- Hild, M. (1998b). Auto-epistemology and updating. *Philosophical Studies*, 92, 321–361.
- Jeffrey, R. (1988). Conditioning, kinematics and exchangeability. In B. Skyrms & W. L. Harper (Eds.), *Causation, chance, and credence*. Dordrecht: Kluwer.
- Kolodny, N. (2005). Why be rational?. *Mind*, 114, 509–563.
- Lewis, D. (1979). Attitudes de dicto and de se. *The Philosophical Review*, 88, 513–543.
- Loewer, B., & Belzer, M. (1983). Dyadic deontic detachment. *Synthese*, 54, 295–318.
- Maher, P. (1992). Diachronic rationality. *Philosophy of Science*, 59, 120–141.
- Plantinga, A. (1993). *Warrant: The current debate*. New York: Oxford University Press.
- van Fraassen, B. C. (1984). Belief and the will. *Journal of Philosophy*, 81, 235–256.
- van Fraassen, B. C. (1985). Empiricism in the philosophy of science. In P. M. Curchland & C. A. Hooker (Eds.), *Images of science*. Chicago: University of Chicago Press.
- van Fraassen, B. C. (1993). Rationality does not require conditionalization. In E. Ullmann-Margalit (Ed.), *The Israel colloquium studies in the history, philosophy, and sociology of science* (Vol. 5). Dordrecht: Kluwer.
- van Fraassen, B. C. (1995). Belief and the problem of Ulysses and the sirens. *Philosophical Studies*, 77(1), 7–37.
- Weisberg, J. (2007). Conditionalization, reflection, and self-knowledge. *Philosophical Studies*, 135(2), 179–197.