

- Schroeder, M. 2008. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Clarendon Press.
- Unwin, N. 1999. Quasi-realism, negation, and the Frege–Geach problem. *The Philosophical Quarterly* 49: 338–52.
- Wedgwood, R. 2007. *The Nature of Normativity*. Oxford: Clarendon Press.

Getting Noncognitivism out of the Woods

By MARK SCHROEDER

In the preface to *Being For* I defend the book's preoccupation with what might seem to many philosophers to be matters of detail. The details, after all, are where the devil is, and accountability to details is what makes the difference between a subject and a discipline. It isn't my goal in the book to focus on details to the exclusion of the big picture; on the contrary I try hard throughout the book to stay focused on details whose resolution is central to the larger questions that are at stake, and to explain their centrality. Still, the book is detail oriented, and it pursues a particular line of inquiry. The main idea, as I explained in the *Synopsis*, is to constructively develop what I take to be the most promising sort of expressivist theory, in order to illustrate the significant sorts of problems that it faces.

One significant consequence of the book's preoccupation with the details arising along this particular line of inquiry is that the book starts fairly far along in the story, leaving a fair number of upstream issues largely undiscussed. This is in part because the topic of the book, metaethical *expressivism*, is just one strand of the noncognitivist tradition stretching, at least, from Hägerström, through Ogden and Richards, Schlick and Carnap, Ayer and Stevenson, Hare, Edwards, Smart, Blackburn, Gibbard and Horgan and Timmons. And it is partly because I have discussed some of those other issues elsewhere – particularly in 'What is the Frege-Geach Problem', in 'Hybrid Expressivism: Virtues and Vices', and in my forthcoming *Noncognitivism in Ethics*. But it is also because the approaches to the 'negation problem' that I discuss and criticize in Chapter 3 are only the ones I take to have the greatest promise. *Being For* includes no discussion at all of an earlier generation of approaches to the Frege-Geach Problem which I've argued elsewhere face a quite different set of problems; for purposes of the book, I merely presuppose that these earlier approaches are unsatisfactory.

On the other hand, it follows from the dialectical feature that the constructive work in *Being For* is ultimately for the sake of a take-down, and that it is natural for critics to look for a place to get off of the boat – and the earlier, the better. This dialectical feature and the undiscussed upstream issues reach a confluence in Andrew Alwood's and Ralph Wedgwood's rich and

illuminating contributions to this symposium. Wedgwood suggests getting off of the boat before my solution to the negation problem even gets started, by offering an alternative solution of his own, which manifests what I take to be an earlier generation of approaches to the Frege-Geach Problem. Alwood is so anxious to get off of the boat he never even gets on – his suggestion is that expressivism is not, after all, the most promising bloom on the noncognitivist tree, and that we would do better focusing our attentions on illocutionary theories more closely akin to Hare’s prescriptivism.

Their contributions therefore give me an opportunity to step back and review some of the reasons why I *do* think that expressivism is particularly worth evaluating, and why I think it *was* right to set aside the earlier generation of expressivist attempts to solve the Frege-Geach Problem. Although they are well intentioned, I’ll be explaining in what follows why I take the principal moves suggested by Wedgwood and Alwood to be problematic, as well as trying to address a few other comments that they make along the way.

1. *Alwood on Mood and Optimism for Illocutionary Theories*

Andrew Alwood’s main concern, in his contribution, is that I’ve short-changed the prospects for illocutionary theories in the noncognitivist tradition – which would be alternatives to expressivism, because rather than explaining the meaning of normative sentences in terms of the nature of normative thought, they would explain the meaning of normative sentences in terms of the nature of normative *assertion*. Hare’s prescriptivism is such an illocutionary theory, and Alwood is much more optimistic about their prospects than about the prospects of expressivism, so though he agrees with me that the prospects of expressivism are dim, he thinks that its prospects were the wrong thing to investigate in the first place – if I really cared about the prospects of views in the tradition including Ayer, Stevenson, Hare, Blackburn and Gibbard, he thinks, I should have focused my energies on illocutionary theories like Hare’s.

One piece of evidence in favour of this conclusion cited by Alwood is that the very arguments that I consider in Chapter 1 for ‘licence for optimism’ in the prospects for expressivism work by paying attention to complex imperative sentences – the ‘licence for optimism’ argument suggests that since we all ultimately need to understand how complex imperative sentences work, we should be optimistic about being able to understand how complex normative sentences work, according to noncognitivism. But Alwood naturally wonders: why should this make us optimistic, unless our noncognitivist theory says that normative sentences *work like imperatives*? Only in that case should we get real licence for optimism from this argument, and so consequently arguments like this should lead us to focus on illocutionary

theories which compare normative sentences to imperatives, rather than on expressivism, which makes no such comparison.

This is fair enough, but I discussed Hare's argument from licence for optimism only because it was the earliest version of such an argument in the literature and the easiest to discuss; in fact Allan Gibbard offers a precisely analogous argument for licence for optimism which stands to expressivism in exactly the same way that Hare's argument stands to licence for optimism for his prescriptivism, by paying attention to complex intentions – negative intentions, conjunctive intentions, disjunctive intentions and conditional intentions. Since everyone needs to be able to understand these complex intentions and how they are related, that is supposed to give us licence for optimism that we should be able to understand complex normative thoughts, if normative thoughts are like intentions.

On the flip side, Alwood also holds that I've short-changed the *extent* to which we can find licence for optimism about illocutionary embedding, by considering too restricted a diet of data; he argues that we should in fact be much more optimistic, particularly if we consider data like his example from Korean, in which imperatives are embeddable in indirect discourse, and cases of 'mixed-mood' conjunctions using connectives like 'but' and 'for'. I agree – the data about the embeddability of mood are particularly striking and particularly complicated, and they present a fertile ground for productive research.

In addition to these considerations about which sort of theory – expressivism or illocutionary theories – are better motivated by arguments from licence for optimism, Alwood argues that expressivist semantics – including as developed in *Being For* – is subject to a general failure to be able to account for the semantics of mood. His primary complaint is that I write throughout of 'ordinary descriptive sentences' all expressing 'belief' and of 'normative sentences' expressing 'states of being for' – where it is clear that what I mean is that 'ordinary *declarative* descriptive sentences' and 'normative *declarative* sentences' do so. Very cursory reflection reveals that I was ignoring non-declarative sentences – interrogatives, imperatives and optatives, for example – entirely. Alwood believes both that this is a serious obstacle to the viability of expressivism, and that a semantic framework that has the right structure to deal adequately with mood (and hence a broadly illocutionary theory of meaning more generally) will itself have the tools to provide an illuminating account of normative language.

On this charge I plead guilty – I *did* ignore the issue of the semantics of non-declaratives throughout *Being For*. From one perspective, of course, it is grist for my mill that mood is one more construction for which expressivists have trouble providing a compositional account. But it is worth noting as well that provisionally ignoring non-declarative sentences for the sake of getting one's theory off of the ground does not distinguish expressivism in any way from the vast majority of mainstream semantic theorizing, in which

mood is largely an afterthought. So in this respect, it is not clear that expressivists are any worse off than any of mainstream, broadly truth-conditional, semantic theorizing.

Alwood speculates (quite prematurely, I think) about how a biforcated attitude semanticist might seek to account for mood, and finds his own speculative proposal wanting; I would be inclined, in contrast, to pay more attention to how mainstream semanticists have sought to account for mood, and to see how expressivists might build on that. In fact, in contrast to expressivism more generally, biforcated attitude semantics has the right structure in order to co-opt quite a lot of what mainstream semanticists have said about mood. For example, according to one common view about the semantics of questions, the semantic value of a question is not a single proposition, but a set of propositions – the set of possible answers to the question. Well, as I discuss in Chapter 11 of *Being For*, in biforcated attitude semantics there is a natural candidate for the role of propositions – they are the pairs of properties that I somewhat neutrally refer to as ‘semantic values’ throughout the book. Biforcated attitude semantics assigns every closed declarative sentence to one of these pairs, just as a more conventional semantic theory assigns every closed declarative sentence to a proposition. So just as a more conventional semantic theory can assign a closed interrogative sentence to a set of propositions, biforcated attitude semantics can go on to assign interrogative sentences to sets of such pairs. I take the ready availability of such a strategy to be part of the cumulative evidence that the solution to the Frege-Geach Problem offered in *Being For* really is on the right track.

Now, I’m all in favour of recognizing mood as an important part of meaning, and I’m fairly sympathetic to the idea that understanding mood will require us to generalize on the semantic resources that have been used to understand the meanings of declarative sentences. I’m particularly intrigued, as I noted above, by the phenomena associated with ‘mixed-mood’ sentences, and think that they are likely to pose a central phenomenon for any adequate theory of mood and hence of meaning in general to ultimately be able to explain. It may even be that focusing on mood gives us the best reasons to be optimistic about the prospects for noncognitivism. But I *don’t* think turning to illocutionary theories is a viable *alternative* to solve the problems that face expressivism.

The reason why not is simple. Everyone needs an account not only of what ‘stealing is wrong’ means, but of what it is to think that stealing is wrong. Expressivists roll these two tasks into one, by proposing to give an answer to the former *by* giving an answer to the latter. But even if we don’t do that, we still, ultimately, need to give an answer to the latter. And that account must ultimately explain what it is to think that stealing is not wrong as some function of what it is to think that stealing is wrong – in a way that allows for an explanation of why it is rationally inconsistent to have both thoughts

at the same time. That is exactly what expressivists have been trying to do, and exactly the task that I try to evaluate and pursue in *Being For*. Turning to an illocutionary-based theory of meaning isn't, therefore, a way of getting out of solving these problems – it is merely committing to solving them *twice* – once for language, and again for thought.

2. A Hybrid Illocutionary Theory?

In concluding his contribution, Alwood goes on to note his belief that the best sort of noncognitivist theory in metaethics will be not only illocutionary based, but hybrid – in his words, that 'normative words have two types of meaning: one ordinary, descriptive meaning and one nondescriptive, illocutionary meaning'. Now, this is a strange development, given Alwood's insistence on the embeddability of imperatives as the ground for our licence for optimism. Hybrid explanations of embeddability turn on terms having a descriptive component, but unless Alwood thinks that imperatives have a descriptive component, his arguments for licence for optimism based on imperatives ought to raise the possibility of an illocutionary view that is *not* hybrid.

Still, as Alwood notes, I've written fairly extensively about hybrid theories elsewhere.¹ And in fact one of the important things I've argued elsewhere is that a hybrid theory helps significantly with solving the Frege-Geach Problem only if it assumes that every sentence with the same moral predicate plays exactly the same expressive/illocutionary role. To see the heart of why this is so, suppose that Alwood is willing to accept both of the premises of the following argument:

1. Lying is wrong.
2. If lying is wrong, then getting little brother to lie is wrong.
- C. Getting little brother to lie is wrong.

Now, accepting 1 and 2 should commit Alwood to going on to accept C (or else giving up at least one of 1 or 2). This is what I call the *inference-licensing property* of valid arguments. But since according to a hybrid view, accepting C involves more than simply accepting some descriptive content – it also involves performing or being willing to perform some illocutionary act – to explain this datum, a hybrid theory must explain why accepting 1 and 2 commits Alwood to *both* accepting the descriptive component of C *and* performing the illocutionary act associated with C. Hybrid views modelled on pejoratives and racial slurs explain this by holding that accepting the descriptive contents of 1 and 2 commits Alwood to accept the descriptive content

1 Schroeder (2009).

of C, and holding that the illocutionary component of C is the *very same as* the illocutionary component of 1 and 2. In contrast, it is clear from Alwood's discussion of the Frege-Geach Problem that he does *not* believe that embedded moral predicates always play the same illocutionary role, so his hybrid view is going to face an important problem, here, and cannot rely on a comparison to more successful hybrid theories.

These are among the reasons I'm fairly sceptical about the direction Alwood proposes to take us; I discuss obstacles facing hybrid theories in substantial detail in Schroeder (2009) and in Chapter 10 of my forthcoming *Noncognitivism in Ethics*. It is unfortunate that I wasn't able to pursue all of these questions in *Being For*, but I firmly believe that it is better to do one thing at a time, and do it well, than to do many things but only poorly. Still, despite these reasons for doubt, I look forward to seeing Alwood pursue his illocutionary theory in the level of detail that I've pursued the development of bifurcated attitude semantics – for that is ultimately the stage at which we will be able to compare their final prospects.

3. *Against Wedgwood*

The bulk of Wedgwood's contribution is built around an alternative expressivist semantic proposal that he offers – not to defend it as ultimately correct, for, like me, Wedgwood is no expressivist – but to argue that it fares at least as well as bifurcated attitude semantics. Whereas the basic idea in *Being For* is that there is a single kind of attitude expressed by *every* sentence in the language, the basic idea of Wedgwood's 'Gibbard-inspired' proposal assumes that there is a single kind of attitude expressed by all *logically complex* sentences, but atomic sentences can express different kinds of attitudes. The way that his compositional rules work is that logically complex sentences express attitudes *towards* the attitudes expressed by simpler sentences, or towards combinations of such attitudes.

By my classification, this makes Wedgwood's proposal a version of a *Higher-Order Attitudes* approach to the Frege-Geach Problem, which makes it strikingly similar to the approach in Chapter 5 of Simon Blackburn's (1984) *Spreading the Word* – an approach that Blackburn himself shortly gave up. So though Wedgwood calls his approach 'Gibbard-inspired' and Wedgwood uses the trick of hyperplanners in order to establish that his approach gets the right results about propositional logic, which sounds like Gibbard's view (Gibbard 2003), I believe that it is to the problems for Higher-Order Attitudes theories that we must turn, in order to see where Wedgwood's approach goes astray.²

2 Blackburn's (1984) approach in *Spreading the Word* is piecemeal, and only introduces an account of the conditional, with no discussion of how that compares to accounts of negation or conjunction, with no verification that the results it generates are right for

Although I don't discuss Higher-Order Attitudes theories in *Being For*, but skip directly to more recent views which I take to be more promising, I do discuss them in 'What is the Frege-Geach Problem?' and in Chapter 6 of my forthcoming *Noncognitivism in Ethics*. Though there are several serious-looking general problems facing Higher-Order Attitudes theories, the one I take to be the most central and illuminating is what I call the *van Roojen Problem*, because it is due to Mark van Roojen's 1996 article, 'Expressivism and Irrationality'. What van Roojen pointed out was that Higher-Order Attitudes theories make the attitudes expressed by logical compounds *too much like* the attitudes expressed by atomic sentences, and as a consequence, they turn out to predict incompatibilities between sentences that are not, in fact, incompatible.

Wedgwood's 'Gibbard-inspired' account in fact makes the van Roojen Problem quite easy to illustrate – nearly as easy as for Blackburn's original Higher-Order Attitudes view. Compare what Wedgwood says about the attitudes expressed by 'ought' sentences and by negations:

'Or consider the judgement that you ought to turn left now; in this Gibbard-inspired theory, this normative judgement can be identified with the attitude of *disagreeing* with the action of not turning left now [– that is, AGAINST(not turning left now)].' (122)

'Take a sentence A that expresses a mental state α . Then we can say that A's negation ' $\sim A$ ' expresses the state of disagreeing with α – that is, AGAINST(α).' (123)

It is striking that 'ought' sentences and 'it is not the case that' sentences express such similar attitudes – and we can exploit this fact in order to construct, for each negated sentence, a non-negated sentence with which it is manifestly not equivalent, but with which Wedgwood's 'Gibbard-inspired' account predicts it to be equivalent.

For example, suppose that we start with Wedgwood's example, 'you ought to turn left now', which according to Wedgwood expresses the attitude, AGAINST(not turning left now). By Wedgwood's account of negation, 'it is not the case that you ought to turn left now' therefore expresses the attitude, AGAINST(AGAINST(not turning left now)). But the same rule that tells us that 'you ought to turn left now' expresses AGAINST(not turning left now), also tells us that 'You ought not to be against not turning left now' expresses AGAINST(not not AGAINST(not turning left now)). But now compare these two attitudes: AGAINST(AGAINST(not turning left now)) and AGAINST(not not AGAINST(not turning left now)). They may not be *exactly* the same attitude (depending on how fine-grainedly Wedgwood individuates the objects of

arbitrary complex sentences. It's helpful to think of Wedgwood's invocation of hyperplanners as a shortcut to settle these matters.

being against), but Wedgwood does need to assume that the same hyperplanners are in each.³ And from that it follows that the sentences which express each are equally inconsistent with ‘you ought to turn left now’. But intuitively ‘it is not the case that you ought to turn left now’ and ‘you ought not to be against not turning left now’ are *not* inconsistent. So Wedgwood’s account succeeds at explaining the inconsistency of ‘you ought to turn left now’ and ‘it is not the case that you ought to turn left now’ only by overgeneralizing.

Since Higher-Order Attitude theories like Wedgwood’s proposal predict the inconsistency of some pairs of sentences that are not actually inconsistent, they also validate some arguments that are not actually valid. For example, compare the following two arguments:

1. You ought to turn left now.
 2. If you ought to turn left now, then you ought to turn on the blinker now.
 - C. You ought to turn on the blinker now.
1. You ought to turn left now.
 - 2*. You ought to either not be against not turning left now or not be against being against turning on the blinker now
 - C. You ought to turn on the blinker now.

On Wedgwood’s account, 2 and 2* express attitudes that are identical up to the interdefinability of ‘and’ and ‘or’, and hence which would be shared by any hyperplanner – consequently he predicts that they are equivalent, and hence that the first argument is valid only if the second is.⁴ Since the second argument does not appear to be valid, this is evidence that not much of an explanation of the validity of the first argument has been given, either.

What the van Roojen Problem shows, I think, is that Higher-Order Attitudes accounts are not formally adequate – because the very same attitude that is expressed by logical compounds is also expressed (or at least expressible) by a predicate, they validate arguments that are not valid.⁵

- 3 To see why, note that you are in the state, AGAINST(not turning left now), just in case you are in the state, not NOT AGAINST(not turning left now). Since these two states are equivalent, they must allow the same hyperplanners. And for a state to allow a hyperplanner is for the hyperplanner to not be against that state. Hence, a hyperplanner must be against one of these states just in case she is against the other – so the same hyperplanners are in each of AGAINST(AGAINST(not turning left now)) and AGAINST(not not AGAINST(not turning left now)).
- 4 For comparison, P2 expresses AGAINST(AGAINST(not turning left now)&AGAINST(AGAINST(not turning on the blinker now))) and P2* expresses AGAINST(not either not AGAINST(not turning left now) or not AGAINST(AGAINST(not turning on the blinker now))). These two states are shared by exactly the same hyperplanners for exactly the same reasons laid out in note 3 above.
- 5 Note that if we follow Gibbard in allowing that the only constraint on which attitudes can be expressed by predicates is that they must simply be attitudes with which it is possible to *disagree*, we *have* to allow that this attitude can be expressed by a predicate.

On the other hand, the main virtue of Higher-Order Attitudes theories is that they clearly provide a *constructive* answer to the question of what mental states are expressed by arbitrary complex sentences. In this, they contrast with the more recent class of approaches to the Frege-Geach Problem that I discuss in Chapter 3 of *Being For*, which are formally adequate but I argue are non-constructive. What I argue in ‘What is the Frege-Geach Problem’ and especially in Chapters 6 and 7 of *Noncognitivism in Ethics* is that this is an essential trade-off, so long as we assume that each normative predicate expresses an unanalysable attitude towards its subject – for example, that ‘stealing is wrong’ expresses an unanalysable attitude towards stealing. Although I took the problems for the earlier generation of Higher-Order Attitudes approaches for granted in *Being For*, the central idea of the book is that it is necessary and sufficient to escape this problematic trade-off, to assume that the attitudes associated with normative predicates are *structured* – just as ‘grass is green’ expresses not a primitive *believes-green* attitude towards grass but the *belief* that grass *is green*, ‘stealing is wrong’ must express not a primitive *disapproval* attitude towards stealing, but *being for* some relation towards stealing. I take this to be a major piece of progress over the Higher-Order Attitude approach as well as over the theories that I discuss in Chapter 3 of the book. This is why I see Wedgwood’s proposal as a step back for expressivism, rather than as a step forward.^{6,7}

4. *Why Biforcated?*

One problem posed by Wedgwood that I take particularly seriously is the question of the point of biforcated attitudes. The problem is not the point, from a technical point of view, of postulating such things; the problem is why we would have states of mind that consist in *two* states of being for, one of which is redundant. If proceeding as if p entails not proceeding as if $\sim p$, so that if I am motivated to proceed as if p , then I am *ipso facto* motivated to proceed as if $\sim p$, then if I am for proceeding as if p , what good does it do me to also be for proceeding as if $\sim p$? Why isn’t being for proceeding as if p

6 For further discussion, see especially Chapters 6 and 7 of *Noncognitivism in Ethics*.

7 There are other reasons why I think it is unpromising to try to explain negation in terms of a higher order attitude of being against. One salient one is that this approach does not generalize to other attitudes with negated contents – for example, to want for stealing not to be wrong, wondering whether stealing is not wrong, or assuming for the sake of argument that stealing is not wrong. Since someone who wonders whether stealing is not wrong does not disagree with someone who wonders whether stealing is wrong, for example, wondering whether stealing is not wrong is not a matter of being against wondering if stealing is wrong – so the Higher-Order Attitudes approach to negation does not generalize in any straightforward way to wondering. In contrast, biforcated attitude semantics assigns ‘stealing is not wrong’ to a semantic value that can function as the object of the attitude of wondering, just as well as it can function as the object of belief.

enough? I've been bothered by this question for some time, and Robert Brandom has also posed me a very similar question in conversation.

I don't really have a very good answer to this question in the form that biforcated attitude semantics actually takes in the book; the best answer I have is to re-characterize 'biforcated attitudes' as states of being for the stronger member of a pair of properties, rather than strictly speaking as being for both members. This is a minimal move (albeit one that unfortunately undermines the name for the view), which preserves all of the structural features of biforcated attitude semantics elaborated in the book. On this conception, we justify the fact that the semantic values of (declarative) sentences are pairs of properties not on the grounds that the attitudes expressed by sentences are pairs of properties, but simply on the ground that the weaker member of the pair keeps track of which attitude is expressed by the sentence's negation.

This perspective is really no different from that taken by a variety of other semantic approaches which share with biforcated attitude semantics the structure of the strong Kleene truth tables – for example, Scott Soames's (1999) paracomplete treatment of truth in *Understanding Truth* or Graham Priest's (1987) dual paraconsistent approach. Each of these theories assigns to each predicate both an extension and an anti-extension, which are assumed in Soames's treatment to be exclusive and assumed in Priest's treatment to be exhaustive, and the role of the anti-extension is to track the behaviour of negations. Biforcated attitude semantics is essentially a non-extensional analogue of Soames's paracomplete theory, and hence it is useful to think about the motivation for assigning sentences to pairs in company with the motivations of such related theories.⁸

Another reason to find this is a useful comparison is that my ultimate complaint about biforcated attitude semantics in *Being For* – that it runs into trouble in extending itself in at least the most obvious way to the treatment of binary quantifiers like 'most' and 'many' – should in principle have an analogue for other strong Kleene-based theories. Though I don't have space to evaluate this question here, this may mean either that the binary quantifiers problem applies to a wider range of views, or it may give us licence for optimism to think that the problem with binary quantifiers will be subject to a general solution.

5. Conclusion

It's been a privilege to have the opportunity here to explore some of the issues both upstream and downstream from the line of argument in *Being For*, and

8 For some reasons to think that biforcated attitude semantics' non-extensional features provide an advantage in dealing with truth in particular, see my 'How to be an Expressivist About Truth'.

to be pushed in particular on issues that I wasn't able to sufficiently explore in the book. Despite the reasons I've articulated here for advising proponents of views in the noncognitivist tradition to steer clear of the 'Woods and their suggestions, I'm still keen to understand how their ideas might be improved. Ultimately, the way we gain better understanding in philosophy is by pushing possible ideas to their limits; *Being For* is an exercise in this conception of philosophical progress, but much more remains to be done.⁹

University of Southern California
3709 Trousdale Parkway
Los Angeles, CA 90089-0451, USA
maschroe@usc.edu

References

- Alwood, A. 2010. Imperative clauses and the Frege–Geach problem. *Analysis* 70: 105–17.
- Blackburn, S. 1984. *Spreading the Word*. Oxford: Oxford University Press.
- Gibbard, A. 2003. *Thinking How to Live*. Cambridge: Harvard University Press.
- Priest, G. 1987. *In Contradiction: A Study of the Transconsistent*. Oxford: Oxford University Press.
- Schroeder, M. 2008a. What is the Frege–Geach problem? *Philosophy Compass* 3/4: 703–20.
- Schroeder, M. 2008b. *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Schroeder, M. 2009. Hybrid expressivism: virtues and vices. *Ethics* 119: 257–309.
- Schroeder, M. forthcoming. How to be an expressivist about truth. Forthcoming in *New Waves in Truth*, eds N. J. Peterson and C. Wright. Basingstoke and New York: Palgrave Macmillan.
- Soames, S. 1999. *Understanding Truth*. Oxford: Oxford University Press.
- van Roojen, M. 1996. Expressivism and irrationality. *Philosophical Review* 105: 311–35.
- Wedgwood, R. 2010. Schroeder on expressivism: For – or against? *Analysis* 70: 117–29.

9 Special thanks to Anthony Ellis, Andrew Alwood, Ralph Wedgwood, Allan Gibbard and Robert Brandom.