



Normative disagreement: a functional account for inferentialists

Sebastian Köhler¹

Published online: 27 March 2020
© The Author(s) 2020

Abstract There was a time when meta-ethical expressivism seemed to be the only game in town for meta-ethical non-representationalists. In recent years, though, meta-ethical inferentialism has emerged as a promising non-representationalist alternative. So far, however, inferentialists lack something that would really allow them to draw level with expressivists. This is an explanation for the distinctive difference between normative and descriptive vocabulary when it comes to disagreement—something expressivists can explain in terms of the difference between representational and desire-like states and which constitutes one of the primary motivations for expressivism. This paper develops a novel and distinctive account for this difference on behalf of inferentialists, based on the different *functions* of these two vocabularies. Not only does this account help inferentialists, it also shows how non-representationalist accounts can capture the relevant disagreement phenomena *without* appealing to the sorts of desire-like states expressivists tend to appeal to.

Keywords Expressivism · Inferentialism · Disagreement · Meta-ethics · Function

1 Introduction

Expressivism has long (relative to meta-ethics' history) held the throne of non-representational meta-ethical views. Recently, though, Matthew Chrisman (2016) has argued that meta-ethical *inferentialism* should hold that position, because it has expressivism's advantages, without its disadvantages. However, one strong

✉ Sebastian Köhler
se.koehler@web.de

¹ Department of Philosophy and Law, Frankfurt School of Finance & Management, 60322 Frankfurt, Germany

argument for expressivism is notably absent from Chrisman's discussion. This is that expressivism can explain a distinctive difference between descriptive and normative vocabulary.

Roughly speaking, for *descriptive* vocabulary our judgements about whether speakers disagree are sensitive to certain divergences in their application of that vocabulary. But, for normative vocabulary our judgements about disagreement are *not* sensitive to similar kinds of divergences. Expressivists explain this through the difference between disagreement in representational belief and disagreement in desire-like states, while representationalists seem unable to properly explain that difference. This argument—call it the “Argument from Disagreement”—is one of the *primary* motivations for expressivism. And so, if inferentialists do not have their own account of the phenomenon it highlights, it is just plain false that inferentialism is superior to expressivism. Hence, the absence of such an account is a *serious* theoretical lacuna for inferentialists.

This paper's aim is to make progress on this unexplored issue. It argues that inferentialists have a response to the Argument from Disagreement that is both novel and distinct from expressivists' by appealing to normative vocabulary's *function*. The idea is, roughly, that if inferentialists tie the way vocabularies admit to disagreement to their respective functions, they can explain the difference between descriptive and normative vocabulary. First, because it is descriptive vocabulary's *function* to keep track of and communicate about features of the environment, it should not be *permissive* regarding certain divergences in application. That is, it should *not* permit wide divergences in application before we must conclude that someone is either misusing the vocabulary or means something else with it than we do. But, for inferentialists the *function* of normative vocabulary is to allow speakers to take an explicit stand regarding and to discuss the inferential roles of sentences in their language. To play this function, though, normative vocabulary must be permissive when it comes to divergences in application.

This account removes inferentialists' theoretical lacuna. And, it marks important meta-ethical progress by showing that non-representationalists can deal with the phenomena highlighted by the Argument from Disagreement *without* invoking desire-like states. This broadens the theoretical resources for dealing with the Argument from Disagreement and shows that non-expressivist non-representationalist views are to be taken very seriously.

The paper proceeds as follows: Sect. 2 presents representationalism, expressivism, and, particularly, inferentialism in greater detail. Section 3 presents the Argument from Disagreement. Section 4 develops the functional response. Section 5 concludes.

2 Representationalism, expressivism, inferentialism

Representationalism, expressivism, and inferentialism are views about the nature of normative thought and discourse. According to *representationalism*, we account for normative thought and discourse in terms of what part of reality they represent. Such views have well-known problems. Non-naturalist representationalists face

metaphysical and epistemological worries, while naturalist representationalists seem to miss what is distinctive about normative thought and discourse.

According to expressivists, we avoid these problems by abandoning representationalism. According to *expressivism*, we account for normative thought and discourse by holding that normative judgements do not play a representational, but a desire-like role. Naturally, expressivism itself has problems, most prominently with fitting its commitments about normative language into a plausible general view about language (exemplified by the Frege-Geach Problem).

Recently, Matthew Chrisman (e.g. 2016) has argued that meta-ethical inferentialism—another non-representationalist view—can take advantage of the arguments supporting expressivism, without its problems (more recently Christine Tiefensee (ms.) has joined the meta-ethical inferentialist camp to argue that meta-ethical inferentialism is the best option for so-called “relaxed realists”). Before we turn to *meta-ethical* inferentialism (“m-inferentialism” in what follows), though, a closer look at inferentialism (“inferentialism” in what follows) more generally will be helpful.

Inferentialism is a view about language more generally and has most extensively been developed in the contemporary debate by Robert Brandom (e.g. 1994, 2000). M-inferentialists tend to build their accounts on his ideas, so here the core ideas of that account will be presented, while trying to avoid some of the complexities and specific theoretical commitments of Brandom’s account. The core idea of Brandom’s account is this:

A performance deserves to count as having the significance of an assertion only in the context of a set of social practices with the structure of (in Sellars’s phrase) a *game of giving and asking for reasons*. Assertions are essentially performances that can both serve as and stand in need of reasons. (Brandom 2000: 189)

What does this mean? The way this has been fleshed out is as follows: On an inferentialist account, an utterance of a declarative sentence, i.e. an assertion, has its content in virtue of its inferential role. Specifically, an assertion has its content, primarily, in virtue of the inferences one is *committed* and *entitled* to drawing by making that assertion. “Inference” here means “material inference,” which is to be understood broadly. *Good* material inferences in this sense are, e.g.

- (1) It is raining, so the streets are wet.
- (2) The match is struck, so it will light.
- (3) Grass is green, so grass is coloured.

Both “commitment” and “entitlement” are technical notions held to capture certain social-normative statuses that are explanatorily basic. The basic idea behind these statuses is that assertions have their content because there is a certain social practice (the “game of giving and asking for reasons”) in which one is committed and entitled to certain things by making it, most importantly drawing certain inferences.

Commitments, in this context, are supposed to be statuses comparable to obligations. By making assertions, speakers undertake commitments to drawing certain inferences. For example, someone who asserts

(4) Grass is green.

undertakes commitment to drawing the inference to

(5) Grass is coloured.

and various other inferences. *Entitlements*, on the other hand, are normative statuses comparable to permissions or justifications, which are important for understanding commitments. This is so, because someone who undertakes a commitment is open to a challenge by other members of the social practice to show entitlement to that commitment, e.g. by backing it up with other commitment that it follows from. In fact, someone who undertakes a commitment *entitles* others to hold her responsible for that commitment. For example, someone who asserts

(4) Grass is green.

might be asked to back to up entitlement to the commitment undertaken by (4), e.g. by pointing out that one has seen many instances of grass and they were green. So, the inferences one is committed to drawing by making an assertion are not just from the statement in question, but also *to* that statement from other commitments that (might) grant entitlement to it.

In short: assertions have their content in virtue of the commitment and entitlement structure in which they stand to other assertions. Let us call a commitment that is a commitment to drawing the sorts of material inferences characteristic for an assertion (whatever those might be) a “doxastic commitment.” Note that because commitments and entitlements are deontological statuses, what inferences one is committed to is a *normative* question. Furthermore, note that inferentialists’ picture is *holistic*: what one is committed to depends on one’s other commitments (Brandom, 1994: 477–482 and Chrisman 2016: 188/189).

As stated above, it is primarily in virtue of the *inferential* connections of the commitments we undertake with them that assertions have their content. However, inferentialists also think that we can undertake commitment to *non-inferential* responses by making assertions. Specifically, we can undertake commitments about the non-linguistic world and how to act within it—inputs and outputs to our linguistic practice can be a part of the commitments we undertake in that practice. This allows us to further distinguish two types of commitments: *empirical* and *practical* commitments. *Empirical commitments* are doxastic commitments that are also connected non-inferentially to states of affairs—they are commitments partially characterized by input conditions. *Practical commitments*, on the other hand, are connected non-inferentially to actions—they are commitments partially characterized by output conditions. Note that for inferentialists, practical commitments are not doxastic commitments. In what follows I will use “inferential role” broadly, to encompass both inferences and these relevant kinds of non-inferential responses.

Because this paper concerns *disagreement*, we also need to address when two speakers *disagree* according to inferentialism. Here we return to inferentialism's holism. To avoid that this leads people to always talk past each other, Brandom (1994: 477–487) offers two models for shared content that allow speakers to disagree.

First, inferentialists can deny that commitments' total inferential role determines content, rather than a relevant subset. This distinguishes what one is committed to by using a sentence in virtue of its meaning (its "semantic inferential role") and other things one is committed to using that sentence (its "non-semantic inferential role").¹ On this view, a sentence used by different speakers means the same thing, if they undertake commitments characterized by the same (or relevantly similar) semantic inferential role. Any additional non-semantic inferential role of their commitments is irrelevant. So, two speakers disagree when one asserts and the other denies e.g.

(6) There are chips in the cupboard.

when both undertake commitments characterized by the same (or relevantly similar) semantic inferential role when asserting (6).

Of course, since inferentialists introduce inferential roles to *explain* meanings, it would be problematic if they had to use meaning to explain what parts of inferential roles are relevant for meaning. Hence, this kind of account requires some way of drawing the distinction between semantic and non-semantic inferential role that does not conflict with inferentialism's explanatory approach. I will later suggest how functions could help with this, but bracket those worries for now and turn to the second view.

On this view, contents are *perspectival*: they can only be specified from a point of view, one that brings in additional commitments undertaken by a speaker. That is, the commitments in virtue of which an assertion has its content yield that particular content in virtue of getting at that content from the speaker's perspective. Information about these contents can be transferred, though, and communication can happen, if speakers "navigate and traverse differences in points of view, to specify contents from different points of view" (Brandom 1994: 485). While this remark might appear a bit cryptic, all it seems to come down to is a specific kind of interpretation:

[E]ven in smooth untroubled cases of communication, if you want to understand what I say, you have to be able to associate with it a sentence that in your mouth expresses the claim that the sentence I uttered expresses in mine. For your understanding it (your knowing what I have committed myself to) involves your being able to trace out the inferences that claim is involved in, the evidential significance of what I have said, in order to know what I am committing myself to. This means knowing how it could function as evidence for you, as well as for me, what claims its endorsement would preclude you, as well as me, from being entitled to, and so on. (Brandom 1994: 510)

¹ This might be akin to the analytic/synthetic distinction, but nothing I say presupposes this.

How does this suggestion deal with disagreement? For this we also need the relation of *incompatibility*, which holds between two commitments, if commitment to the one precludes entitlement to the other. With this notion in place, a plausible suggestion is that two speakers Anna and Ben disagree regarding commitment C, if Anna undertakes C and Ben could not undertake C after navigating and traversing differences in points of view without this precluding him from being entitled to one of his other commitments. Hence, if they were to engage in mutual interpretation, they'd realize that their commitments are *incompatible*.

Which of these accounts should inferentialists use? Fortunately, my purposes do not require a decision: the notion of a function will help either model with the Argument from Disagreement. So, let's return to meta-ethics.

What is m-inferentialists' view of normative thought and language? Brandom's (1994: 229–271), Chrisman's (2016: 193–201), and Tiefensee's (ms.) proposals share a common core. This is that normative vocabulary is a species of "meta-conceptual" vocabulary [this idea comes from Sellars (1956, 1958)].² *Meta-conceptual vocabulary* differs from other vocabularies in that it serves to make inferential roles explicit that would otherwise be implicit: it allows us to explicitly articulate and discuss inferential roles. For Brandom all kinds of vocabularies fall into this category, e.g. logical and semantic vocabulary.

To get clearer on what meta-conceptual vocabulary does, consider Brandom's m-inferentialism, on which normative claims make explicit good practical inferences. Recall the distinction between doxastic commitments—those we undertake through assertions—and practical commitments—those non-inferentially connected to actions. Importantly, doxastic and practical commitments can stand in inferential relations to each other and when they do so, the doxastic commitments are supposed to provide entitlement for the practical commitments. So, there are inferences in which practical commitments follow from doxastic ones, such as (Brandom 1994: 245),

(7) I am a bank employee going to work, so I shall wear a necktie.

When someone engages in inferences of these kind—call them "practical inferences"—this is *practical reasoning*. Such reasoning is *good*, if it makes inferences that one is actually entitled to.

On Brandom's m-inferentialism, normative claims serve to explicitly endorse practical inferences as having entitlement. So, a normative sentence such as

(7*) Bank employees ought to wear neckties.

expresses the speaker's commitment that there is entitlement to draw certain practical inferences (here inferences like those in (7)). That and whether certain doxastic commitments offer entitlement for certain practical commitments is what normative vocabulary articulates. Specifically, normative claims allow us to

² Brandom calls this sort of vocabulary "expressive," but in the meta-ethical dialectic we should avoid this label, given that it suggests that m-inferentialism is a form of meta-ethical expressivism. This is certainly *not* the case: m-inferentialists' meta-conceptual vocabulary is very different from what meta-ethical expressivists take "expressive" vocabulary to do.

explicitly undertake commitments in a way that opens the question whether some practical inference is one that one is entitled to draw, to challenge and discussion.

For all m-inferentialists normative vocabulary is a species of meta-conceptual vocabulary.³ They disagree, though, on what kind of meta-conceptual vocabulary it is, e.g. on what kinds of inferences it is supposed to make explicit (for different views to Brandom's, see Chrisman 2016: 179–181 and Tiefensee ms.). However, these differences won't matter here. What matters is only that m-inferentialists hold normative vocabulary to be meta-conceptual vocabulary. I take this to be m-inferentialism's *core* commitment.

So, what speaks for m-inferentialism? M-inferentialism's most prominent and vocal proponent is Chrisman, who argues that m-inferentialism shares expressivism's benefits, while avoiding its problems. For example, he thinks (2011 and 2016: 191–196 and 223–229) that it evades non-naturalism's worries, because it characterizes normative commitments through ontologically conservative inferential roles. And, according to Chrisman (2016: 217–222), m-inferentialism fits into a naturalistic world-view and captures plausible connections between normative judgements and action. Unlike expressivism, though, it is easily embedded into a general account of linguistic content, namely inferentialism (Chrisman 2011, 2012, 2016). Furthermore, unlike expressivism, it is not committed to normative judgements being directive, desire-like states, a claim Chrisman (2016: 180/181 and 197–200) argues is implausible.

However, even if we concede all of this, there is (at least) one issue where expressivism has a crucial edge. This is an account of how to deal with the central difference between descriptive and normative vocabulary when it comes to disagreement. This brings us to the argument that will be this paper's focus.

3 The argument from disagreement

The Argument from Disagreement is very familiar in meta-ethics, so I will only provide a brief summary.⁴ The argument's starting point is the distinction between genuine disagreement and merely talking past each other. Take, again

(6) There are chips in the cupboard.

Assume Anna asserts (6) and Ben denies it. Are they disagreeing? They might not be, if Anna speaks British English and Ben American English. In this case, they would mean different kinds of fried potatoes by "chips." However, they would probably be genuinely disagreeing, if they came from the same linguistic

³ The assumption that normative vocabulary is a species of meta-conceptual vocabulary might strike some readers as intuitively implausible. Here I am not going to defend this assumption, given that my investigation concerns whether it can do the required explanatory work.

⁴ Janice Dowell (2016) questions the probative value of the intuitions employed in the Argument from Disagreement. I bracket this important challenge here, because what concerns me is whether inferentialism is supported by the argument, *if* it works.

community, because then they'd mean the same things with "chips." I take it that this phenomenon is familiar and widely recognized.

This phenomenon is relevant for meta-ethics, because it matters for meaning and content: genuine disagreement seems to presuppose shared meanings and content. Hence, when speakers are talking past each other and when they are genuinely disagreeing, is relevant for the individuation of meanings and concepts. So, any plausible view about the nature of normative terms and concepts must capture this phenomenon for normative terms and concepts and must have plausible implications about when speakers are talking past each other and when they are genuinely disagreeing.

The puzzle begins once we recognize that our intuitions about whether people are disagreeing or talking past each other are distinctively different for normative terms and concepts than they are for paradigmatically descriptive ones. Take the example of "chips" again. Anna and Ben use the same string of letters in their sentences. However, if they speak different idioms of English, we will find that they apply those strings to very different things. This makes us inclined to think that "chips" does not have the same meaning in Anna and Ben's mouth.

For normative terms and concepts, however, such divergence in application does not seem to matter, rather than, broadly speaking, what people do when they use those terms and concepts. To illustrate, take R.M. Hare's well-known example, in which a missionary meets a tribe of cannibals. The missionary does not know the cannibals' language. However, she knows they have a word that like "good" is "the most general adjective of commendation." By using this word, the missionary seems able to communicate and disagree with the cannibals about normative matters, *as long as she uses the term to commend*, despite the fact that she and the cannibals apply that word in radically diverging ways. After all, the missionary applies it to "people who are meek and gentle and do not collect large quantities of scalps [whereas the cannibals apply it to] people who are bold and burly and collect more scalps than the average" (Hare 1952: 148). This indicates that the term's role to commend determines its content, rather than certain patterns of application. This is distinctively different from descriptive terms, where divergence of application would make us inclined to think that they mean different things.

Of course, this is not true for all descriptive terms. At least for natural kind terms, radical divergence in application is compatible with sameness of content. However, normative terms seem distinctively different even from these kinds of terms, as Terence Horgan and Mark Timmons (e.g. 1992) highlight with their comparison of Putnam's (1975) Twin Earth case with a variant for normative terms. What this comparison suggests is that for natural kind terms (and other terms that function similarly), but not normative terms, there is some relevant relation R between *our* talk and some relevant thing (e.g. "water" and H₂O), such that any talk not standing in that relation to that thing cannot be in disagreement with our talk. For normative terms, unlike for natural kind terms, consequently, it is not any relation in which the relevant talk stands to things which seems relevant for content, but its role.

What explains this difference? Note that what is required to deal with this challenge is an account of what disagreement consists in (or an account *in virtue of* *what* statements disagree) that can account for the distinctive differences when it

comes to disagreement between descriptive and normative terms.⁵ On first sight, representationalists seem unable to offer such an account.⁶ For them, normative terms work just like descriptive terms, so they lack the resources to deal with such dis-analogies. Expressivists, on the other hand, can explain the difference. Their explanation turns on the difference between disagreement in representational states and disagreement in desire-like states. The idea is, roughly, that disagreement in representational states cannot survive the relevant divergences, while disagreement in desire-like states can, something for which expressivists have offered various models [e.g. disagreement in attitude (Stevenson 1937), disagreement in plan (Gibbard 2003: 68–75), or disagreement in prescriptions (Ridge 2014: 185–191)]. Hence, if we take descriptive sentences to express representational states, but normative sentences to express non-representational motivational states, we can explain the difference between descriptive and normative language, when it comes to disagreement.

What about m-inferentialists? Can they explain the phenomena in question? So far, m-inferentialists have not offered any account that deals with this challenge. Note that neither of the accounts of disagreement presented in Sect. 2 combined with m-inferentialists' views about normative concepts is sufficient to deal with this challenge. In both cases we would still lack a crucial component for explaining why there is the relevant difference between normative and descriptive terms. After all, telling us that normative vocabulary is meta-conceptual vocabulary does not tell us in virtue of what normative vocabulary has a semantic inferential role that would account for its difference to descriptive vocabulary when it comes to disagreement. Similarly, telling us that normative vocabulary is meta-conceptual vocabulary does not tell us in virtue of what speakers' would be inclined to interpret use of normative vocabulary in a way that would account for the difference between normative and descriptive vocabulary when it comes to disagreement.

Of course, on both accounts one can simply *assume* that normative vocabulary differs in the relevant respect from descriptive vocabulary—but this is simply assuming what the argument challenges us to explain. The absence of such an explanation, however, is a serious theoretical lacuna, given that the Argument from Disagreement is one of the primary motivations for expressivism (see e.g. Gibbard 1992, 2003; Hare 1952, 1981; Horgan and Timmons 1992; Ridge 2014; Stevenson 1937). Given the prominence of this argument in supporting expressivism, the absence of a feasible m-inferentialist explanation makes it highly doubtful that m-inferentialism is the better non-representational view. Hence, m-inferentialists should, urgently, address this issue.

Note also, that they should not just address the issue: they need an account that is different from expressivists'. After all, if m-inferentialists offered an account that committed them to the idea that normative judgements belong to the desire-like

⁵ Thanks to an anonymous referee for *Philosophical Studies* for suggesting that the details of the challenge be put precisely on the table at this point.

⁶ Of course, representationalists *have* responded to this worry (e.g. Björnsson and Finlay 2010; Dunaway and McPherson 2016; Plunkett and Sundell 2013, or van Roojen 2006). The discussion of such responses goes beyond this paper's purposes.

states expressivists appeal to, this would suggest that m-inferentialism is not actually a distinctive meta-ethical position. Rather, it would be a *version* of expressivism located within inferentialism. Hence, to preserve the idea that m-inferentialism is a distinctive *meta-ethical* position m-inferentialists need to offer an account that is substantially different from expressivists’.

I suggest that drawing on theoretical resources already implicit in inferentialism does the trick: this is an appeal to functions. Specifically, it is the idea that certain features of the inferential role of use of vocabularies are, at least in part, explained by appeal to their *function*. In particular, I suggest that on an inferentialist picture, we should think that the commitments we undertake by using certain vocabularies are, at least in part, characterized by their particular inferential roles, *because* of their distinctive function. This idea helps m-inferentialists with the problems raised by the Argument of Disagreement, because it allows them to connect the distinctive difference regarding disagreement between normative and descriptive vocabularies to their respective functions. The next section develops this response.

4 The functional turn: inferentialism and disagreement

This section proceeds as follows: It starts in Sect. 4.1 by briefly explaining how the notion of a function is to be understood here and what theoretical role it is supposed to play. It then motivates that functions so understood fit into inferentialism. Section 4.2 then fleshes out how functions in this theoretical role provide m-inferentialists with a response to the Argument from Disagreement. Specifically, for each of the two inferentialist models for disagreement introduced in Sect. 2, it explains how functions can help m-inferentialists with the Argument from Disagreement.

4.1 Inferentialism and functions

As I said, the account rests on the idea that inferentialists should think that the commitments we undertake by using certain vocabularies are, at least in part, characterized by their particular inferential roles *because* of their distinctive function. To understand this idea, let’s start with how the notion of a function is to be understood here.⁷

Broadly speaking, as this notion is to be understood here, the function of a feature x is to F , if F -ing is what explains why x exists (see Wright, 1973). Applied to linguistic items such as vocabularies, we might say that the function of a vocabulary is what explains why the vocabulary is used in the relevant social practice. Of course, there are many different ways how this broad understanding of the notion of a function could be fleshed out in this context (see e.g. Bedau 1992; Godfrey-Smith 1994; Kitcher 1993; Millikan 1984, 2005 or Neander 1991). But, the purposes of

⁷ The following section has profited significantly from the constructive comments of a referee for *Philosophical Studies* who has pressed me on the question how the suggested account is supposed to be understood. Thanks to the referee for their encouraging and incredibly helpful feedback.

this paper do not require any more robust commitment either way, so I will mostly continue to rely on an intuitive understanding of it. The aim of the paper is to argue that m-inferentialists have a distinctive response to the Argument from Disagreement if they use the notion of a function. And for this it suffices to only present the structure of the explanations available to m-inferentialists to account for the difference between descriptive and normative disagreement—a structure that actually might fit different more concrete views of functions available to inferentialists.⁸

So, what is the theoretical role that this notion is supposed to play on the account that will be suggested? In principle, we could assign at least two theoretical roles to this notion within an inferentialist framework. First, we might say that it only explains why it is that we have vocabularies with particular sorts of inferential roles. Here, we'd assume that all features of vocabularies' inferential roles can be explained *independently* of the particular functions of those vocabularies, and that functions only explain why those vocabularies exist. Secondly, we might say something more robust, namely that vocabularies have their particular inferential roles (at least partially) *in virtue of* the functions they play. That is, particular features of the inferential role of a vocabulary are constitutively explained by the function of these vocabularies. On this reading it is (at least partially) constitutive of e.g. descriptive vocabulary's inferential role that it enables the vocabulary to perform a particular function. Note that on this second approach, hence, the function of a vocabulary at least partially constitutes the meaning of that vocabulary, via its constitutive role vis-à-vis its inferential role.⁹

It is this second interpretation of functions' theoretical role that will be adopted here. For only this second interpretation can do the relevant explanatory work regarding the Argument from Disagreement. Specifically, the way the accounts I will suggest should be read is that uses of different vocabularies disagree in the way they do *in virtue of* the functions these different vocabularies play. I will flesh out shortly what this means on the two inferentialist models of disagreement laid out earlier—though let me flag here already that on these two models the *specific* explanatory roles of functions will differ, leading to two different inferentialist accounts of how functions impact inferential roles. But let me first motivate that the general idea actually fits into inferentialism.

At first sight, one might think that the suggestion that function is partially constitutive of meaning departs from inferentialism's explanatory approach. After all, it is a core idea of inferentialism that it is *inferential role* that explains content. This scepticism, however, would be too quick. First, note that my suggestion does

⁸ My understanding of the notion of a function here follows so-called etiological accounts of that notion. However, inferentialists' explanatory purposes might also be served by a "dispositional" notion of a function, where functions explain what something *contributes* to the capacities of a system (e.g. Bigelow and Pargetter 1987 or Boorse 1976). Thanks to Peter Schulte for helpful discussion and feedback on this issue.

⁹ The idea that functions can play constitutive explanations of meaning is not new, of course. It plays a prominent role in teleofunctionalist accounts [most prominently defended by Ruth Millikan (e.g. 1984, 2005)]. However, in an inferentialist framework this idea is given a novel spin, because here functions play constitutive explanations for inferential roles and in that way impact meaning.

not conflict with the core inferentialist ideas: it is compatible, for example, with the idea that content is to be explained in terms of commitment and entitlement, which are still explanatorily basic notions. After all, even if we assign functions the role I suggest, there is still a sense in which inferential role explains content. It is just that for at least some central aspects of meaning inferential role does not play the most fundamental explanatory role. Instead it is the *function* of the relevant vocabularies that does so. But, the function plays this role *by* having a certain explanatory role vis-à-vis inferential role. So, functions are partially constitutive of meaning only because they are partially constitutive of inferential roles.

Second, my suggestion seems already implicit in existing inferentialist accounts. For example, both Huw Price (e.g. 2011 and 2013: 34) and Michael Williams (2010, 2013) seem to rely on a similar suggestion in their inferentialist accounts. Indeed, it seems to me that we also find it in Brandom's work. For example, Brandom's (1994: 247–249) account of the meaning of normative vocabulary does not at all appeal to inferential role, but only to its *function* as meta-conceptual vocabulary (this is true for any of the other m-inferentialist's too. See Chrisman 2016: 193–200 and Tiefensee ms.). The same holds for Brandom's account of the meaning of any other meta-conceptual vocabulary (e.g. 1994: 107–109). This suggests that for meta-conceptual vocabularies, their *function* as meta-conceptual vocabulary has explanatory priority on Brandom's account and that the vocabularies make their contributions to inferential role (at least partially) in virtue of their specific meta-conceptual function. Similarly, when Brandom (1994: 221–223) explains why empirical vocabulary comes with a default entitlement (which is part of its inferential role), his explanation appeals to the fact that empirical vocabulary has the *function* of being a regress stopper. So, again, this suggests that empirical vocabulary makes its contributions to inferential role (at least partially) in virtue of its specific regress-stopping function.

I think these remarks should suffice to motivate that the notion of a function has a plausible role to play in inferentialism when it comes to constitutively explaining certain parts of the meaning of vocabularies.¹⁰ I will now present how m-inferentialists can use this notion to account for the difference between normative and descriptive vocabulary when it comes to disagreement.

4.2 Putting functions to use

Let me start by considering what the functions of descriptive and normative vocabulary are. Consider first paradigmatically descriptive vocabulary. Here I will use an account that is inspired by (though by no means identical to) Brandom's account of empirical commitments. As mentioned earlier, some commitments are

¹⁰ Of course, I am not saying that every inferentialist has to endorse my suggestion or that my reading of these authors has to be correct. Maybe they are deeply sceptical of it. In this case, we should see my account as a rival inferentialist account, one that offers m-inferentialists a way to deal with the Argument from Disagreement that other inferentialist accounts would lack. Here, though, I will proceed on the assumption that my suggestions fit into existing inferentialist accounts, rather than posing an alternative to them.

not only characterized by inferential connections to other commitments, but also by non-inferential input conditions. Specifically, some commitments are undertaken by responding to features of one's environment. What is special about (some) such commitments is that they come with default entitlement, if they are undertaken under certain kinds of conditions that guarantee the speaker's reliability. Further above I called these "empirical commitments." My suggestion is to think of descriptive vocabulary as that vocabulary used to undertake such commitments.¹¹ And, given that we are entitled to such commitments just in case they are formed reliably in response to the environment, the function of such vocabulary—what explains why a vocabulary allows us to undertake commitments of this kind—is, clearly, to enable us to keep track of and communicate information about features of our environment.

What about normative vocabulary? M-inferentialism's *core* commitment is that it is a species of meta-conceptual vocabulary. This serves the function of making certain sorts of inferences explicit that characterize our commitments. Specifically, the function of normative vocabulary is to allow speakers to take an explicit stand on the inferential connections characteristic of terms in their language in a way that makes the goodness of those inferential connections available for further discussion and deliberation.

Let us now consider how these ideas might be used to account for the phenomena highlighted by the Argument from Disagreement. As I said above, I will proceed by considering, in turn, the two inferentialist models for disagreement introduced above and explain how, on each model, an appeal to functions allows m-inferentialists to deliver the required explanation for the difference between normative and descriptive vocabulary when it comes to disagreement.

I start with the first model. On this view, one denies that the total inferential role of a commitment undertaken by an assertion is responsible for its content, rather than just a relevantly individuated subset, the *semantic* inferential role. On this account, for any declarative sentence S, two speakers talk past each other when one utters S and the other its negation, if the commitments they undertake by use of S are characterized by (relevantly) different semantic inferential roles. On the other hand, they would be disagreeing, if these commitments were characterized by the same (or a relevantly similar) semantic inferential role.

In terms of this model the Argument from Disagreement poses the challenge to explain why certain differences in the use of e.g. "chips" between linguistic communities entail that the commitments undertaken by speakers when asserting

(6) There are chips in the cupboard.

are characterized by different semantic inferential roles, while this is not the case for similar differences in uses of "ought" in e.g.

(8) One ought to give at least 10% of one's income to charity.

¹¹ Note that this is an illustration how the account would proceed, not a final account.

So, the challenge is to explain, on the one hand, why the fact that speakers of British English apply “chips” to different kinds of fried potato than speakers of American English makes a difference to the semantic inferential role of the commitment these respective speakers undertake by using (6). And, on the other hand, why the fact that e.g. speakers of one community use “ought” to apply to actions that maximize utility and speakers of another community use “ought” to apply to actions that satisfy the Categorical Imperative, does not entail differences in the semantic inferential role of the commitment these speakers undertake by using (8). Hence, specifically, the challenge is to explain why descriptive vocabulary makes a contribution to the semantic inferential role of sentences in which it figures that substantially limits what commitments one can undertake by using it, while normative vocabulary does not. With the functional approach, we have an explanation.

First, functions provide a systematic and non-circular way to distinguish contributions to *semantic*, from contributions to *non-semantic* inferential role. Specifically, we can say that the contribution a vocabulary makes to the semantic inferential role of commitments undertaken by employing that vocabulary is restricted to what is necessary for that vocabulary to perform its distinctive function. All other contributions it makes to the inferential role of commitments are only part of the non-semantic inferential role of those commitments. If this is successful, this allows m-inferentialists to draw the distinction between semantic and non-semantic inferential role in a way not in conflict with inferentialism’s explanatory approach. But, more importantly, this way of cashing out the difference between semantic and non-semantic inferential role also accounts for the difference between our two respective vocabularies when it comes disagreement.

Start with descriptive vocabulary. Descriptive vocabulary serves the function of enabling speakers to keep track of and communicate about features of their environment. For descriptive vocabulary to perform this distinctive function, speakers who use it must, via such use, undertake commitments characterized by suitable input conditions: “chips” can only be used to keep track of our commitment about chips, if the commitments undertaken by using “chips” are to a sufficient degree responsive to the presence of chips. Specifically, for descriptive vocabulary to play this function someone who undertakes such a commitment must be committed to her being in a set of conditions in which she is relevantly responsive to the feature in question. These conditions can, of course, vary to some degree with context, depending on how hard it is to keep track of the relevant feature and so on. But, there still will have to be a relatively fixed and stable set of conditions for normal situations associated with the descriptive term in question.¹² Otherwise, it is simply questionable whether and how the use of the vocabulary keeps track of and allows communication about features in our environment.

So, for descriptive vocabulary to perform its distinctive function certain sorts of input conditions must characterize the commitments speakers undertake by use of

¹² At the very least, how speakers undertake or retract these commitments must display suitable deference to experts.

the vocabulary. Specifically, they must be characterized by input conditions that plausibly grant entitlement to the commitment, when satisfied. But, note that these conditions must be such that in general speakers situated in relevantly similar ways should all undertake a commitment characterized by the input in question—at least if they are suitably motivated. Again, otherwise it would be questionable that the use of the vocabulary keeps track of and allows communication about features in our environment.

On the view I suggested here, the use of descriptive vocabulary, therefore, contributes such input conditions to the semantic inferential role of commitments undertaken by its use, because these input conditions are crucial for descriptive vocabulary to perform its distinctive function. This means, though, that we can explain why competent speakers cannot use descriptive vocabulary in (certain) radically diverging ways: the commitments one can undertake by using such vocabulary are constrained by the input conditions characteristic of the semantic inferential role of such commitments.

Hence, if we encounter a different linguistic community which has a term such as e.g. “chips,” but its speakers apply that term (sufficiently) differently to us, we must conclude that their “chips” means something different than ours does. This is so, because our “chips” functions to keep track of certain features in our environment and hence, constrains the range of commitments speakers can undertake by its use. If they apply “chips” (sufficiently) differently, we must conclude that it does not have the same function and, hence, is not characterized by the same contribution to semantic inferential role. In this way, we can explain why certain differences in the application of descriptive terms mean that speakers are talking past each other.

What about normative vocabulary? Normative vocabulary is *meta-conceptual* vocabulary. Specifically, its function is to allow speakers to take an explicit stand on the inferential connections characteristic of terms in their language in a way that makes the goodness of those inferential connections available for further discussion and deliberation. But, to fulfil this function, normative vocabulary’s contribution to semantic inferential role must be significantly thin. For example, when I say

(7*) Bank employees ought to wear neckties.

what I should centrally be committed to in virtue of the meaning of (5*), really, is just that inferences like

(7) I am a bank employee going to work, so I shall wear a necktie.

are *good*, as well to drawing such inferences in suitable conditions. However, the contribution “ought” makes to the semantic inferential role of commitments undertaken by its use should impose *no* (or very little) restrictions on what material inferences I can undertake commitment to as being good. Otherwise, “ought” could not fulfil its function as a piece of meta-conceptual vocabulary. Hence, the contribution “ought” makes to the semantic inferential role of commitments undertaken by its use imposes no (or very little) restrictions on divergences in patterns of use of “ought.” Rather, it only imposes restrictions on the *structure* characteristic for commitments undertaken by its use. Specifically, all commitments

undertaken by use of “ought” will have the structure characteristic of commitments to the effect of certain inferences being good.

Of course, if this is the case, then, when we encounter a different linguistic community who has a term “ought,” but whose members apply that term differently than we do, this need not lead us to conclude that their “ought” means something different than ours does. We should only conclude this when we discover that their “ought” does not serve the function of making certain inferential connections explicit, e.g. when we discover that the commitments they undertake by the use of “ought” are not characterized by the sort of structure that characterizes the commitment we undertake by use of that term.

In this way, m-inferentialists can offer a distinctive take on the suggestion that it is not what things speakers apply normative terms to, but what they *do* with those terms that matters for whether speakers have a genuine normative disagreement. For m-inferentialists, however, this has nothing to do with the kinds of desire-like states expressivists appeal to. Instead, it has something to do with normative vocabulary’s meta-conceptual function. Note that this is even true on Brandom’s account on which normative vocabulary makes explicit good practical inferences: what’s doing the work in accounting for disagreement is not an appeal to directive mental states. Instead it is normative vocabulary’s *meta-conceptual function*. So, on the model that employs the distinction between semantic and non-semantic inferential role, functions can explain the difference between descriptive and normative vocabulary in a way that distinctively differs from expressivists’ explanation.

Let’s turn to the second inferentialist model for how to explain disagreement. This was the view that contents are *perspectival* in the sense that they can only be specified from a specific point of view and that speakers need to engage in mutual interpretation to see what another speaker means. On this view, communication requires the ability to not only know what others are committed to, but what oneself would be committed to, if one used the relevant sentence. On this view, two speakers are talking past each other when one asserts and the other denies e.g.

(8) One ought to give at least 10% of one’s income to charity.

if they would discover, when they “navigate and traverse differences in their points of view” that the expressed commitments are not incompatible. They disagree, if such navigation would reveal that those commitments are incompatible.

Put in terms of this model, the challenge posed by the Argument from Disagreement is to explain, on the one hand, why e.g. the fact that speakers of British English apply “chips” to different kinds of fried potato than speakers of American English makes those speakers, when they interpret each other, inclined to hold that they do not undertake incompatible commitments when the former speakers assert and the latter deny

(6) There are chips in the cupboard.

And, on the other hand, it needs to be explained why comparable differences in application would be irrelevant to speakers when it comes to determining whether

they are undertaking incompatible commitments when one asserts and the other denies a normative sentence like

- (8) One ought to give at least 10% of one's income to charity.

More specifically, on this model, the challenge is to explain why speakers would be more *tolerant* of stark differences in application of normative vocabulary when judging whether the commitments undertaken by use of a normative sentence and its negation are incompatible, than they would be for descriptive vocabularies. The functional approach provides an explanation.

Specifically, my suggestion is that on this model, we should see considerations about the *function* of vocabularies as anchors that guide speakers' interpretations of each other. When confronted with someone uttering a sentence, speakers will attempt to figure out what commitment that speaker has undertaken and what they themselves would be committed to if they used that sentence. One thing that speakers will try to figure out here is whether the commitment the speaker undertakes is *incompatible* with any of *their* own commitments. I suggest that in determining at least this issue, speakers are implicitly governed by considerations about the *function* of the vocabulary in question—they are implicitly guided by the assumption that commitments are incompatible or not, at least in part, in virtue of their respective functions.¹³

When confronted with someone using what looks like paradigmatic descriptive vocabulary, e.g.

- (6) There are chips in the cupboard.

they will implicitly be guided by the assumption that it has the function of enabling speakers to keep track of and communicate about features of their environment. Because they attribute this function, speakers will not be tolerant towards relevant divergence in the application of the term "chips" by a community of speakers in the following sense: If they found a community of speakers that exemplified such a divergence in the use of "chips," they would *not* be interpreting those speaker when they use (6) as undertaking commitments that are incompatible with the commitment *they* themselves would undertake when they negate (6). The explanation for this proceeds along similar lines to the explanation further above: descriptive vocabulary serves the function of reliably communicating information about the environment and, hence, must be characterized by certain sorts of inputs. If these inputs diverge too strongly from one's own in situations in which one is similarly situated, one must conclude that the commitments the speaker undertakes using that vocabulary are not incompatible with one's own.

On the other hand, if the speaker employs what they take to be *normative* vocabulary, e.g. when uttering

¹³ It is important to read this in a way that is not overly intellectualized. I take it that for Brandom, the process of interpretation is a form of know-how—a practical ability that allows speakers to gain knowledge of the commitments others undertake. The account offered here should be understood as a rational reconstruction of this ability, i.e. as a model that enhances our understanding e.g. of what factors it would be sensitive to.

- (8) One ought to give at least 10% of one's income to charity.

they will (implicitly) assume that it has the function to allow speakers to take an explicit stand on the inferential connections characteristic of terms in her language in a way that makes the goodness of those inferential connections available for further discussion and deliberation. Because they attribute this function, though, they will tend to tolerate a divergence in application, when it comes to determining whether use of some sentence undertakes a commitment that is incompatible with their own (obviously, speakers need not be tolerant of opposing and divergent normative views *per se*—it's just that they will not take this divergence to matter *in determining incompatibility*). Hence, they'd tolerate a relatively wide divergence in application of normative terms, when it comes to determining whether the commitment undertaken by (8) is incompatible with their own. The reason why speakers will be so tolerant, is because given the function of normative vocabulary and the perspectival nature of contents, such a divergence in application is to be expected and to be worked with. They will only view the other person as *not* undertaking a commitment incompatible with one of their own, if their interpretation lets them conclude that in that person's language sentences like (8) are not meta-conceptual. It seems, consequently, that if we use the second model, an appeal to the function of the respective vocabularies can also account for the difference between normative and descriptive vocabularies when it comes to disagreement.

Hence, it seems that on both models m-inferentialists can use functions to explain the difference between descriptive and normative vocabulary highlighted by the Argument from Disagreement. The explanation m-inferentialists offer as to why *normative* vocabulary behaves fundamentally different to descriptive vocabulary when it comes to disagreement is, thereby, crucially different from expressivists' explanation: it is its *function* as *meta-conceptual vocabulary* that explains why divergence in the application of normative vocabulary is compatible with sameness of content. This explanation does not allude to the desire-like states expressivists invoke in their account of normative judgement and, hence, is distinct and fundamentally different from expressivists' explanation. Furthermore, given that m-inferentialists can offer an explanation for the phenomena that the Argument from Disagreement draws our attention to, this shows us that m-inferentialism is supported by that argument just as much as expressivism is. At least with regards to this argument, m-inferentialism shares expressivism's virtues.

5 Conclusion

This paper has offered a way for m-inferentialists to deal with the Argument from Disagreement. The suggestion was that m-inferentialists can give an account of the distinctive differences between descriptive and normative vocabularies when it comes to disagreement, by appealing to these vocabularies' *functions*. The core idea was for inferentialists to hold that vocabularies admit to disagreement in certain sorts of ways *in virtue of* their respective functions. I fleshed out how this idea helps with accounting for the differences highlighted by the Argument from Disagreement

by developing, for each of the two general inferentialist models for disagreement, an account that lays out the structure of how this m-inferentialist explanation would proceed. On the first kind of model, the explanation would proceed by tying the *semantic* inferential role of commitments to their respective functions. On the second kind of model, the explanation would proceed by tying speaker's verdicts about incompatibility in the process of mutual interpretation to considerations about functions. If what I have argued is correct, then m-inferentialists can use functions on both models to explain the difference between descriptive and normative vocabulary highlighted by the Argument from Disagreement. And in that case, m-inferentialism is at least on a par with expressivism when it comes to explaining these phenomena.

One last thing relevant for the meta-ethical debate is worth drawing attention to. This is that the framework I developed should *in principle* open the door for other non-representationalist responses to the Argument from Disagreement. The m-inferentialist explanation is based on the idea that normative vocabulary has a meta-conceptual function. This explanation works because a vocabulary can only play this function if it is permissive in the sense discussed here, i.e. allows room for divergence in application. However, there is no reason to rule out that other functions that vocabularies could play might also require such room for divergence. Hence, it might be possible to keep the functional framework I've developed here, but have a different view on the function of normative vocabulary and still be able to deal with the phenomena highlighted by the Argument from Disagreement. Of course, the response will not work if we assume that normative and descriptive vocabularies have the same functions. But, other theoretical options might be feasible. In fact, while I developed the account here within an inferentialist framework, similar responses could plausibly be developed within other relevantly similar views in the philosophy of mind and language, such as, for example conceptual role views more generally. Hence, not only does the response I have developed help m-inferentialists. It also moves the debate forward by opening conceptual space for a variety of non-representational alternatives to expressivism.

Acknowledgements Open Access funding provided by Projekt DEAL. For helpful comments and discussion on the paper in its various stages, I would like to thank Emad Atiq, Annika Böddeling, Matthew Chrisman, Lars Dänzer, Daan Evers, Miguel Hoeltje, Nick Laskoswski, Errol Lord, Christian Nimtz, Raphael van Riel, Mike Ridge, Neil Roughley, Hanno Sauer, Peter Schulte, Paulina Sliwa, Christine Tiefensee, audiences at the Oberseminar Philosophische Anthropologie at Duisburg-Essen University, the Oberseminar Theoretische Philosophie at Duisburg-Essen University, and the GAP 10, as well as an anonymous referee for *Philosophical Studies*.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bedau, M. A. (1992). Where's the good in teleology. *Philosophy and Phenomenological Research*, 52, 781–805.
- Bigelow, J., & Pargetter, R. (1987). Functions. *Journal of Philosophy*, 86, 181–196.
- Björnsson, G., & Finlay, S. (2010). Metaethical contextualism defended. *Ethics*, 121(1), 7–36. <https://doi.org/10.1086/656304>.
- Boorse, C. (1976). Wright on functions. *The Philosophical Review*, 85, 70–86.
- Brandom, R. (1994). *Making it explicit. Reasoning, representing, and discursive commitment*. Cambridge MA: Harvard University Press.
- Brandom, R. (2000). *Articulating reasons. An introduction to inferentialism*. Cambridge, MA: Harvard University Press.
- Chrisman, M. (2011). Expressivism, inferentialism and the theory of meaning. In M. Brady (Ed.), *New waves in metaethics* (pp. 103–125). New York: Palgrave MacMillan.
- Chrisman, M. (2012). On the meaning of “Ought”. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 7, pp. 304–332). Oxford: Oxford University Press.
- Chrisman, M. (2016). *The meaning of “Ought”: Beyond descriptivism and expressivism in metaethics*. Oxford: Oxford University Press.
- Dowell, J. (2016). The metaethical insignificance of moral twin earth. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 11, pp. 1–27). Oxford: Oxford University Press.
- Dunaway, B., & McPherson, T. (2016). Reference magnetism as a solution to the moral twin earth problem. *Ergo*, 3(25), 639–679.
- Gibbard, A. (1992). *Wise choices, apt feelings*. Cambridge MA: Harvard University Press.
- Gibbard, A. (2003). *Thinking how to live*. Cambridge, MA: Harvard University Press.
- Godfrey-Smith, P. (1994). A modern history theory of functions. *Nouûs*, 28(3), 344–362.
- Hare, R. M. (1952). *The language of morals*. Oxford: Oxford University Press.
- Hare, R. M. (1981). *Moral thinking: Its levels, method, and point*. Oxford: Oxford University Press.
- Horgan, T., & Timmons, M. (1992). Troubles for new wave moral semantics: The ‘open question argument’ revived. *Philosophical Papers*, 21(3), 153–175. <https://doi.org/10.1080/05568649209506380>.
- Kitcher, P. (1993). Function and design. *Midwest Studies in Philosophy*, 18, 379–397.
- Millikan, R. (1984). *Language, thought, and other biological categories*. Cambridge MA: MIT Press.
- Millikan, R. (2005). *Language: A biological model*. Oxford: Clarendon Press.
- Neander, K. (1991). The teleological notion of ‘function’. *Australasian Journal of Philosophy*, 69(4), 454–468. <https://doi.org/10.1080/00048409112344881>.
- Plunkett, D., & Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, 13(23), 1–37.
- Price, H. (2011). One cheer for representationalism? In H. Price (Ed.), *Naturalism without mirrors*. Oxford: Oxford University Press.
- Price, H. (2013). *Expressivism, pragmatism and representationalism*. Cambridge: Cambridge University Press.
- Putnam, H. (1975). The meaning of “meaning”. *Minnesota Studies in the Philosophy of Science*, 7, 131–193.
- Ridge, M. (2014). *Impassioned belief*. Oxford: Oxford University Press.
- Sellars, W. (1956). Empiricism and the philosophy of mind. In H. Feigl & M. Scriven (Eds.), *Minnesota studies in the philosophy of science* (Vol. 1, pp. 253–329). Minneapolis, MN: University of Minnesota Press.
- Sellars, W. (1958). Counterfactuals, dispositions, and the causal modalities. In H. Feigl, M. Grover, & M. Scriven (Eds.), *Minnesota studies in the philosophy of science* (Vol. 2, pp. 225–308). Minneapolis MN: University of Minnesota Press.
- Stevenson, C. L. (1937). The emotive meaning of ethical terms. *Mind*, 46(181), 14–31. <https://doi.org/10.1093/mind/XLVI.181.14>.
- Tiefensee, C. (ms.). Metasemantics for the Relaxed. provisionally accepted for publication. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics*. Oxford: Oxford University Press.
- van Roojen, M. (2006). Knowing enough to disagree: A new response to the moral twin earth argument. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 1, pp. 161–194). Oxford: Oxford University Press.

- Williams, M. (2010). Pragmatism, minimalism, expressivism. *International Journal of Philosophical Studies*, 18, 317–330.
- Williams, M. (2013). How pragmatists can be local expressivists. In H. Price (Ed.), *Expressivism, pragmatism and representationalism* (pp. 128–144). Cambridge: Cambridge University Press.
- Wright, L. (1973). Functions. *The Philosophical Review*, 82, 139–168.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.