

Quasi-naturalism and the Problem of Alternative Normative Concepts

Camil Golub

The following scenario seems possible: a community uses concepts that play the same role in guiding individual actions and shaping social life as our normative concepts, and yet refer to something else. As Eklund (2017) argues, this apparent possibility poses a problem for any normative realist who aspires to vindicate the thought that reality itself favors our ways of valuing and acting. How can realists make good on this idea, given that anything they might say in support of the privileged status of our normative concepts can be mirrored by the imagined community? For instance, realists might claim that using our concepts is what we *ought to do* if we are to describe normative facts correctly, but people in the other community can truthfully say the same about their concepts, using their own concept of *ought*.

A promising approach to this challenge is to try to rule out the possibility of alternative normative concepts in the first place, by arguing that any concepts that have the same normative role must share a reference as well, at least when it comes to authoritatively normative concepts like *wrong* or the all-things-considered *ought*. Eklund calls this *referential normativity*. But this is only the outline of a solution to the problem. The realist still needs to provide a metasemantic picture that supports referential normativity for the relevant concepts. In this paper I argue that *normative quasi-naturalism* offers such a theory of normative content that avoids Eklund's challenge.

Quasi-naturalism has three main ingredients: (1) expressivism about normative discourse, i.e. the view that normative claims express desire-like mental states; (2) a deflationary account of certain metaphysical and semantic notions in terms of which realism is usually stated, such as truth and fact; and (3) naturalist realism about normativity, i.e. the view that normative properties are identical with, or fully constituted by, objective natural

properties, where such claims about property identity or constitution are understood as internal to normative discourse and amenable to an expressivist treatment.

On this view, if two concepts encode the same desire-like attitudes and have the same social function, then they are the same concept, even if they are otherwise used in vastly different ways. In other words, sameness of normative role entails sameness of concepts. I call this *conceptual normativity*, to be distinguished from referential normativity, especially because the two need not always align: for instance, conceptual normativity arguably holds for etiquette predicates, but such predicates can have a different reference for different communities. Authoritatively normative concepts, however, exhibit both conceptual and referential normativity: for example, “wrong” refers to the same natural property for us and in any other context where it has the same normative role, where this claim about reference is again understood as a normative judgment. Or so the quasi-naturalist will argue.

In developing this response to Eklund's challenge, I will rely on a hybrid version of expressivism according to which normative claims express both desire-like states and representational beliefs, and which fits into a broader contextualist semantics for deontic and evaluative terms (Ridge 2014). This view allows quasi-naturalists to capture the distinctive theses of metaphysical naturalism, and to distinguish authoritatively normative concepts from other normative concepts, for which realism and referential normativity do not hold true.

At the end, I will address some objections to my proposal. More precisely, I will respond to Eklund's worry that even views that adopt referential normativity might not escape the problem of alternative normative concepts, because concepts that have slightly different normative roles can still be in competition in a way that undermines normative realism, and will address the concern that quasi-naturalism is not a form of genuine realism, given its expressivism and its deflationary framework.

1. The problem of alternative normative concepts

Here is the scenario that Eklund (2017) asks us to imagine:

“**Alternative.** There is a linguistic community speaking a language much like English, except for the following differences ... While their words ‘good,’ ‘right,’ and ‘ought’ are associated with the same normative roles as our words ‘good,’ ‘right,’ and ‘ought,’ their words aren’t coextensive with our ‘good,’ ‘right,’ and ‘ought.’ So even if they are exactly right about what is ‘good’ and ‘right’ and what ‘ought’ to be done, in their sense, and they seek to promote and to do what is ‘good’ and ‘right’ and what ‘ought’ to be done in their sense, they do not seek to promote what is good and right and what ought to be done.” (18)

If Alternative is possible, Eklund argues, this poses a problem for any *ardent realist* about normativity, that is, any realist who tries to vindicate the thought that “reality itself favors certain ways of valuing and acting” (2017: 1). This idea seems to be undermined by the alarming symmetry between our concepts and the ones employed by the imagined community: users of those alternative concepts seem to be getting things right just as much as we do. To be sure, they do and say things that are “wrong,” in our sense of the term, but we equally do and say things that are “wrong” in their sense. And any claim we might make about the privileged status of our normative concepts—e.g., that using our concepts is what we ought to do if we are to describe normative reality correctly—can be mirrored by the other community, using their own normative concepts and with equal justification.

This challenge superficially resembles the Moral Twin Earth problem raised by Horgan and Timmons (1991) for certain versions of naturalist realism: it similarly relies on a scenario where some concepts have the same normative role—that is, the same role in guiding practical deliberation and interpersonal advice and criticism—but are applied to different things. However, the two problems differ in at least two important ways. First,

Eklund's challenge does not rely on linguistic intuitions about sameness of meaning or reference, which the realist might try to debunk or explain away (see, e.g., Dowell 2016); on the contrary, it relies on the supposition that the relevant concepts have a different reference. Secondly, the problem of alternative normative concepts affects not only naturalist views that adopt a causal theory of reference for normative terms, but any realist view that allows for concepts with the same normative role and a different reference.¹

To be sure, Cornell-style realist views committed to a causal theory of reference (e.g., Boyd 1988) are the most obvious target for Eklund's challenge. If the reference of a normative term is the natural property, or cluster of natural properties, with which its use is causally connected in the right way, then reference can clearly come apart from normative role: two communities can use terms that play the same practical role in individual deliberation and interpersonal criticism and yet bear relevant causal links to different things in the world.² But again, the problem arises for any realist view on which the reference of normative terms is determined in ways that have little to do with their normative role.

This includes, for instance, Jackson's (1998) analytic descriptivism, on which the reference of moral terms is determined by certain platitudes that form the core of the folk theory of morality: e.g., the reference of "wrong" is whatever makes true certain widely accepted and deeply entrenched beliefs associated with the use of "wrong." As Eklund (2017, 21) points out, if folk platitudes are to play this reference-fixing role, they will have to include substantive moral claims that apply moral terms to specific actions, states of affairs,

¹ Eklund's Alternative scenario also resembles the famous missionary example used by Hare (1952) in his argument against naturalism. But again, Eklund's challenge applies to a wider range of realist views, and is not meant to elicit intuitions about disagreement with respect to such cases.

² Boyd's theory of reference is more complex than this quick summary suggests: in particular, it is causal links *that tend to bring about true predication* which determine the reference of normative terms. (See Väyrynen (2018) for a discussion of how this affects the Moral Twin Earth problem.) However, this qualification does not matter much in the context of Eklund's challenge, as the view still allows for alternative normative concepts.

etc. (Merely structural platitudes specifying how normative terms relate to each other would be insufficient for this purpose.) But again, two terms can arguably have the same normative role while being associated with vastly different sets of such moral claims.

Moreover, non-naturalist versions of realism are also vulnerable to this challenge if they posit mechanisms of reference determination for normative concepts that come apart from the normative role of said concepts. For instance, the view that normative terms ascribe certain non-natural properties because those properties are reference magnets in virtue of their metaphysical eliteness (Dunaway and McPherson 2016) seems to allow for scenarios in which certain terms refer to different properties for different communities, even while playing the same role in practical deliberation and interpersonal criticism. This gives rise, again, to a symmetry between the two communities which should alarm the realist: there seems to be no good way to argue that one community is objectively right and the other is objectively mistaken in their normative practices.

As Eklund puts it, any realist who accepts the possibility of alternative normative concepts faces a choice between two unattractive options: “either accept that there are ineffable questions about normativity, or else accept that the importance of normative questions is deflated” (2017, 19). That is, the first option is to claim that, even if each community is using normative concepts correctly given what those concepts refer to, there is a further question about which set of concepts is genuinely better or privileged by normative reality, the answer to which will favor our concepts. However, Eklund argues, realists will also have to concede that this further question cannot be adequately articulated, let alone settled, using our normative concepts—or any other set of normative concepts, for that matter. In other words, it is an ineffable question. Alternatively, realists might be content to defend the privileged status of our normative concepts using those very concepts: for instance, they might claim that members of the alternative community do and say things that

are objectively wrong, in our sense of “wrong,” and this is sufficient for a defense of realism in the face of Alternative-type scenarios. But this option does not seem to properly vindicate ardent realism either, given that similar arguments will be available to users of alternative normative concepts. We are still left without a clear case for the idea that reality itself favors our ways of valuing and acting, and disfavors other ways of using normative concepts.

Now, some realists will resist Eklund's dilemma. They will accept that there can be alternative normative concepts but will argue that the alarming symmetry between such concepts and ours can be broken in a way that does not take us into the realm of ineffability and is not deflationary. McPherson (2020), for instance, argues that some normative concepts are privileged because they refer to *authoritatively normative properties*, where this notion is understood in metaphysical terms. Eklund is skeptical that this kind of appeal to metaphysics can help realists avoid his dilemma and resolve the problem of alternative normative concepts, and I am inclined to agree with him, but I will not discuss this issue in the present paper. I am interested in exploring the other main strategy for defending realism in the face of Eklund's challenge, which is to deny the very possibility of alternative normative concepts.

2. Referential normativity

The most promising strategy in the face of Eklund's problem is to try to avoid it in the first place, by arguing that there can be no alternative normative concepts.³ This means offering a metasemantic picture on which concepts that have the same normative role have the same reference as well. Eklund uses the label *referential normativity* for this alleged feature of normative concepts.⁴

³ More precisely, no alternative authoritatively normative concepts. More on this, below.

⁴ To be more precise, Eklund suggests that realists should aim to make good on the idea that there can be *non-defectively referentially normative* concepts: that is, concepts whose reference is determined by their normative role, and which are not empty or wildly semantically indeterminate. From here on, I will use *referential*

Now, as McPherson (2020) and Plunkett (2020) point out, merely proposing an account of normative concepts that rules out Alternative-type cases seems insufficient for vindicating ardent realism: even if sameness of normative role entails sameness of reference, this does not in itself ensure that our normative concepts refer to objective features of the world, or that they are favored by normative reality. However, a charitable reading of Eklund is that he takes referential normativity to be a *necessary* condition for vindicating ardent realism in the face of his challenge, which might also be sufficient if other conditions for realism are met, such as genuine objectivity for normative truths.

Here is the gloss that Eklund gives for the notion of normative role, which plays a key role in the idea of referential normativity and in this discussion more generally:

“It is characteristic of normative predicates that they are fit to be used in practical deliberation relating to what to do; it is characteristic of such predicates that their application has, so to speak, practical consequences in addition to merely theoretical ones. They have normative roles (...) [I]t is likewise intuitive that different kinds of normative predicates are used differently, as far as such normative roles are concerned. But if one can compare normative roles, one can also judge when two predicates have the same normative role. Even if one cannot state in very informative terms what the normative role of a predicate is, one can intelligibly speak of sameness of normative role.” (38)

This description of normative role might seem loose, but it has the virtue of leaving open a range of options for capturing this notion. For instance, an expressivist characterization of the practical function of normative concepts will look quite different from a cognitivist account

normativity and *referentially normative* to mean *non-defective referential normativity* and *non-defectively referentially normative*.

of normative role, but both views will aim to capture the same intuitive feature of normative concepts: their role in guiding individual deliberation and interpersonal criticism.⁵

One more preliminary point before introducing quasi-naturalism. Alternative-type scenarios only pose a problem for realism when they involve *certain* normative concepts, like the moral *wrong* or the all-things-considered *ought*. No realist should be troubled by scenarios in which, say, etiquette concepts have the same normative role but a different reference, because such concepts are not plausibly amenable to a realist account in the first place: we *should* expect their reference to vary between communities, even if they have the same normative role across different contexts of use. For this reason, Eklund restricts his challenge to *thin* normative concepts, but the relevant divide here is arguably not between thin and thick concepts, but rather between normative concepts that are *authoritatively normative* and those that are not (McPherson 2020). In any case, this is how I will understand the task of a view that aims to address Eklund's challenge by defending referential normativity: to rule out the possibility of alternative authoritatively normative concepts.

This more precise understanding of the task also carries another important implication, however: the realist needs to provide a metasemantic picture that upholds referential normativity for authoritatively normative concepts, but also allows for the reference of *other* normative concepts to vary across different contexts of use. I will argue that quasi-naturalism can accomplish this twofold task.

⁵ Moreover, various conceptions of normative role might be put to use in defending referential normativity. For example, one theory of normative meaning that would support referential normativity is Wedgwood's (2007) conceptual role semantics, which provides the framework for his version of non-naturalist realism. On this picture, the semantic roles of normative terms consist in certain rules that govern the use of said terms in practical reasoning, and these semantic roles uniquely determine the reference of normative terms: for instance, the basic rule for *ought* is that judging that one ought to ϕ rationally commits one to intending to ϕ , and the reference of *ought* is, roughly speaking, the property that makes such plans *correct* or genuinely choiceworthy (Wedgwood 2007: 80, 104, 153). Once again, however, I will set aside such alternative options for addressing Eklund's challenge, as my focus is on pursuing an expressivist approach to this problem.

3. Introducing quasi-naturalism

Normative quasi-naturalism has three main ingredients: (1) expressivism about normative discourse; (2) deflationism about certain metaphysical and semantic notions in terms of which realism is usually stated, such as truth, fact, or description; and (3) naturalist realism about normativity, i.e., the view that normative properties are identical with, or fully constituted by, objective natural properties. In this section I will provide more details about each of these pieces and will explain how they fit together.

Expressivism is, broadly speaking, the view that normative claims express desire-like mental states, where this is understood as a thesis about normative meaning, and not just about the pragmatics of normative discourse. Here I will rely on a specific version of expressivism, namely a *hybrid* view on which normative claims express both desire-like attitudes and corresponding representational beliefs.⁶ This view is well placed to account for the metaphysical claims of naturalist realism and can fit into a broader contextualist semantics for deontic and evaluative terms, a feature which plays an important role in properly addressing Eklund's challenge.

Before bringing contextualism into the picture, however, let me introduce a simple version of hybrid expressivism (Ridge 2007, Toppinen 2013), which will do for the purposes of showing how expressivists can accommodate metaphysical naturalism about normativity. On this view, an atomic normative sentence like “Cheating on your taxes is wrong” expresses a relational mental state that consists in (1) an attitude of disapproval of actions that have a certain natural property, and (2) a belief that cheating on your taxes has that natural property. Such relational mental states can be realized by different combinations of desire-like attitudes

⁶This is in contrast to *pure expressivism*, according to which the semantic content of normative claims fully consists in certain desire-like mental states, like plans or attitudes of approval and disapproval. See, e.g., Blackburn (1984, 1993), Gibbard (1990, 2003), or Schroeder (2008).

and corresponding beliefs for different speakers. For example, “Cheating on your taxes is wrong,” when uttered by a utilitarian, might express an attitude of disapproving actions that do not maximize utility and the belief that cheating on your taxes fails to maximize utility, while an ethical egoist's utterance of the same sentence might express disapproval of actions that are not in the speaker's best interest and the belief that cheating on one's taxes is not in the speaker's best interest, but both utterances will express the same relational mental state.⁷

Here is how expressivism thus understood is compatible with naturalist realism about normativity. First, as is well known, deflationism about truth, facts, and other related notions allows expressivists to endorse many tenets of realism. For instance, on a deflationary account, truth is not a substantive property and the function of the notion of truth is merely expressive: it allows us to endorse and reject first-order claims in any domain of discourse, and to abbreviate and generalize when talking about claims that we endorse or reject. That is, on this picture, “It is true that p ” is equivalent to p , and this equivalence schema fully captures the meaning of “true”. Thus, expressivists can hold that “It is true that cheating on your taxes is wrong,” taking this claim to simply rehearse the first-order normative judgment that cheating on your taxes is wrong (see, e.g., Blackburn 1984, 1993; Gibbard 2003). More generally, on this account, any judgment is truth-apt if it meets certain minimal conditions concerning its structure and its interaction with other parts of language and thought,⁸ and normative judgments do meet these conditions. Similarly, expressivists can claim that there are normative *facts*, or that normative judgments *describe* such facts, by relying on deflationary accounts on these notions.

⁷ See Schroeder (2013), Toppinen (2013), and Ridge (2014) for different articulations of this idea.

⁸ For instance, according to the version of deflationism known as *disciplined syntacticism*, any sentence that has a declarative form and is governed by sufficient standards of warrant is thereby truth-apt. See Lenman (2007) for a discussion of the relation between expressivism and different versions of disciplined syntacticism.

Moreover, expressivists can endorse claims about *objective* normative truths and facts, by taking such claims to express a particular kind of attitude. For instance, on the simple version of hybrid expressivism articulated above, the claim that “It is an objective fact that cheating on your taxes is wrong” might express a relational mental state comprising (1) an attitude of disapproving actions that have a certain natural property, even when considering scenarios in which we ourselves did not disapprove of such actions; and (2) the belief that cheating on your taxes has that natural property.⁹

Everything said so far about the compatibility between expressivism and realist commitments is familiar quasi-realist fare, and does not depend on a hybrid expressivist framework. When it comes to accommodating the metaphysical claims of naturalist realism, however, hybrid expressivism becomes particularly useful, as Bex-Priestley (forthcoming) has argued. Take the following naturalist thesis:

(A) Moral wrongness is identical with failing to maximize utility.¹⁰

(A) is both normative and metaphysical in nature, and is easily amenable to a hybrid expressivist treatment. To use the same simple version of hybrid expressivism, (A) can be understood as expressing (1) an attitude of disapproving action types that have a certain property, and (2) the belief that the property in question is identical with failing to maximize utility. Similar hybrid expressivist accounts can be given for naturalist theses about

⁹ See Blackburn (1984, 1993) and Gibbard (2003) for similar accounts of objectivity claims, but in a pure expressivist framework.

¹⁰ There is a dispute among naturalists about whether the relation between normative and natural properties is best understood in terms of property identity or in terms of some other relation like constitution or reduction. See, e.g., Brink (1989), who argues against property identity naturalism on the grounds that moral properties are fully realized by certain natural properties, but they *could* have been realized by many other sets of natural properties, and perhaps even by certain supernatural properties. I am setting this issue aside here because it is not relevant for Eklund’s challenge and hybrid expressivism can capture both kinds of metaphysical claims.

constitution relations between normative and natural properties, or other metaphysical relations that fall short of identity.¹¹

This is then how expressivists can endorse metaphysical naturalism about normative properties. But someone might wonder *why* expressivists should be naturalists, especially given that their deflationary approach to metaphysical issues seems to align them with quietist non-naturalist realists such as Dworkin (1996) and Scanlon (2014).

Gibbard (2003) argues that expressivists must be naturalists, because any coherent planner is committed to what he calls the Natural Constitution thesis: there is a natural property that constitutes what is good, or what one ought to do.¹² It is important to note, however, that Gibbard uses “natural” in a broad sense that potentially includes *supernatural* properties such as being willed by God.¹³ Bex-Priestley (forthcoming) offers a similar transcendental argument for moral naturalism, but in a hybrid expressivist framework: the judgement that moral properties are irreducibly normative will always involve a false belief, he argues, because all properties toward which we can have desire-like attitudes can be described in non-normative terms—or at least this is what expressivists should think, given

¹¹ To be clear, I am not claiming that only hybrid expressivism can accommodate the metaphysical claims of normative naturalism. Perhaps pure expressivism can also achieve this, although I am inclined to agree with Bex-Priestley (forthcoming) that it lacks the resources to distinguish between identity or reduction claims, on one hand, and mere correlation claims involving normative and natural properties, on the other: e.g., between “Goodness reduces to being pleasurable” and “Necessarily, things are good if and only if they are pleasurable”. In any case, here I am only making the positive claim that hybrid expressivism does provide a clear account of metaphysical naturalist claims. As mentioned before, I rely on hybrid expressivism for this reason and because it can fit into a broader contextualist semantics for deontic and evaluative terms, a feature which I believe plays an important role in properly addressing Eklund's challenge.

¹² Fundamentally, Gibbard's argument does not even depend on expressivism: anyone committed to the supervenience of the normative on the natural is thereby committed to the Natural Constitution thesis, he argues, and expressivism only plays a role here insofar as it entails supervenience.

¹³ Gibbard also uses the term “prosaically factual” for the category of properties to which normative properties can be reduced. This is arguably a better label since it does not suggest that these must be the kind of properties studied by the natural and social sciences.

that they aim to explain the content of normative claims without relying on irreducible normative concepts. However, as Bex-Priestley himself acknowledges, this argument too only establishes that normative properties are reducible to non-normative or descriptive properties, rather than naturalism in a narrow sense that excludes supernatural properties.

This being said, I believe that naturalism in the narrow sense is the most natural option for expressivists, for several reasons. First, naturalism fits with the original motivations for the expressivist project (Blackburn 1984, Gibbard 1990): the aim to provide an ontologically parsimonious theory of normative discourse, on which fundamentally there are only natural facts and our attitudes directed at those facts, but which also leaves normative discourse in good standing.

Secondly, naturalism arguably has epistemological advantages over non-naturalist realism, given that it takes normative facts and properties to be causally efficacious: for instance, naturalists can better explain the reliability of our normative beliefs, by appealing to causal connections of the right sort between our beliefs and the normative facts.

Thirdly, there is arguably no good reason *not* to be a naturalist about normativity, particularly if one combines naturalism with expressivism. Properly defending this claim would take me beyond the scope of this paper, but let me just say that none of the standard arguments against naturalism, like Moore's (1903) open question argument or Enoch's (2011) argument that normative properties are "just too different" from natural properties for naturalism to be true, seem effective against versions of naturalism that acknowledge important differences between normative and purely descriptive concepts, particularly the fact that normative concepts have a practical function that involves action guidance and social coordination. Such functional differences between normative and purely descriptive concepts can be used to explain away the intuitions behind many of these arguments against

naturalism. And of course, expressivism is the most salient way of articulating the idea that the primary function of normative concepts is practical rather than descriptive.

One more clarificatory point. I use the term *quasi-naturalism* for this view that combines expressivism and normative naturalism not because of its deflationary framework or because I take it to be less than fully realist, but because it stands in contrast with representationalist accounts of normative meaning, such as causal theories of reference for normative terms or other metasemantic pictures associated with naturalism.¹⁴ For quasi-naturalists, normative concepts are primarily individuated through their practical roles rather than through representational relations with the normative realm. This is what allows quasi-naturalists to avoid the challenge of alternative normative concepts, to which I now return.¹⁵

4. Quasi-naturalism as an Alternative-unfriendly view

Quasi-naturalism can vindicate ardent realism in the face of Eklund's challenge by ruling out the possibility of alternative authoritatively normative concepts while allowing that other normative concepts, such as etiquette predicates, can have a different reference for different communities. Here is how this works.

First, we need to distinguish between Conceptual and Referential Normativity:

Conceptual Normativity (CN) Sameness of normative role entails sameness of concepts.

Referential Normativity (RN) Sameness of normative role entails sameness of reference.

¹⁴ Majors (2006) uses the term “quasi-naturalism” slightly differently: for him, “quasi” refers to the fact that moral properties are not recognized as real in virtue of their explanatory role.

¹⁵ Quasi-naturalism resembles in some respects a view that Eklund calls *presentationalism*, and which he takes to be another salient option for addressing the problem of alternative normative concepts. The core idea of presentationalism is that “normativity generally is something that characterizes only our devices for representing the world and not what these devices stand for” (Eklund 2017, 67). However, presentationalism denies the existence of normative properties, while quasi-naturalism claims that normative properties do exist, even if they are identical with, or reducible to, natural properties, and even if their normativity simply consists in being the kind of properties to which normative concepts refer.

If a view delivered (CN), this might seem to be enough to forestall Eklund's challenge, because (CN) seems to make Alternative-type scenarios impossible.¹⁶ Moreover, the simple version of hybrid expressivism stated in the previous section upholds (CN) for the concepts to which it applies: for instance, if the term “wrong” has the same normative role for us and for the imagined community in Alternative, this is in virtue of expressing the relational mental state partly composed of a desire-like mental state, the expressivist will argue, which also entails that both communities use the same concept of wrongness.

However, this appeal to (CN) is too quick and does not properly address Eklund's challenge. This is because it still seems possible that different communities use the same concepts but with a different reference, and that there is no good way to argue that one community's use of concepts is somehow privileged by normative reality. Indeed, as mentioned before, this is precisely the type of scenario that we should expect when it comes to *non-authoritatively* normative concepts such as etiquette predicates. A proper defense of realism in the face of Eklund's challenge needs to distinguish such relativistic concepts, which arguably exhibit (CN) but not (RN), from authoritatively normative concepts like *wrong* and the all-things-considered *ought*, for which both (CN) and (RN) must be upheld.

In order to achieve this goal, I propose that quasi-naturalists rely on a more complex version of hybrid expressivism that fits into a broader contextualist semantics for deontic and evaluative terms. In what follows I will use a specific view that fits this bill, namely Ridge's (2014) version of hybrid expressivism, but this is meant to illustrate a more general point: it is the structure of Ridge's view that does the work here, rather than its details.

The general case for contextualism is by now well-known (e.g., Finlay 2014, Ridge 2014, Chrisman 2016, Worsnip 2019). Deontic and evaluative terms like *ought* and *good* are

¹⁶ In a similar vein, Plunkett (2020, fn. 11, p. 7) suggests that the core issue raised by Alternative-type scenarios concerns normative concepts as such rather than concept extensions or properties.

used in many different ways—for instance, we can talk about what we prudentially ought to do, or about what we ought to do according to the standards of British etiquette, or about what we ought to do, all things considered—and it is implausible that such terms are ambiguous between these different uses. The more plausible hypothesis is that terms like *ought* and *good* have a unique meaning or character but different semantic contents in different contexts of use, in the same way that indexicals like *here* and *I* have a uniform character but context-sensitive contents.

Ridge relies more specifically on a Kratzer-style contextualism for deontic and evaluative terms (Kratzer 2012). On this model, two parameters determine the semantic content of terms like *ought* and *good* in particular contexts of use: (a) a modal base, i.e. a body of propositions held fixed in context, or a set of worlds in which certain background conditions are met; and (b) an ordering source that ranks worlds in the modal base, for instance by how well they satisfy certain standards.

Here is an example: “You ought to give a large part of your income to charity” will be true, in any context of utterance, just in case “You give a large part of your income to charity” is true in all of the worlds in the modal base that are highly ranked according to the ordering source. The ordering source in each context will depend on the flavor of the given deontic or evaluative claim. For instance, in some contexts, the ordering source might consist in norms of etiquette, while in another context worlds will be ranked in accordance with French law, and so on.

For Ridge, authoritatively normative claims are a special subset of uses of deontic and evaluative terms, for which the ordering source consists in acceptable standards of practical reasoning. For example, “You ought to give a large part of your income to charity,” when

used in an authoritatively normative sense, means that any acceptable standard of practical reasoning would recommend that you give a large part of your income to charity.¹⁷

It is important to note that this contextualist semantics does not entail relativism or any other form of anti-realism about normativity. Normative realists can adopt this semantic model and claim that the ordering source in authoritatively normative contexts of use consists in *objectively acceptable* standards of practical reasoning, or some similar parameter, making thus good on the idea that, when we make authoritatively normative claims, we are talking about what we really ought to do, or what is objectively wrong, rather than about what follows from some set of standards that we or others around us happen to accept.¹⁸

Ridge also wants to make good on realist ideas about normative objectivity, but in an expressivist framework. His motivation is that authoritatively normative judgments have certain distinctive features that push us toward expressivism: they are action-guiding, intimately connected with emotions, and they allow for the possibility of radical disagreement. For this reason, Ridge proposes an expressivist account of the notion of acceptable standards: to judge that a standard is acceptable is to endorse a normative perspective that does not rule out that standard, where normative perspectives are understood as noncognitive practical stances. More precisely, standards of practical reasoning are rules

¹⁷ Ridge simply uses the word *normative* for this category of claims, but I am restating his view in terms of authoritative normativity because I think we should allow that domains like etiquette and law are also normative in a weaker sense: they exhibit what is sometimes called *formal* or *generic* normativity (see, e.g., Copp 2005, McPherson 2011, Baker 2017).

I should also note that I have glossed over some important differences between Kratzer's and Ridge's versions of contextualism. In particular, on Ridge's view, the standards that constitute the ordering source for modal terms like "must" and "ought" rank actions rather than possible worlds when such terms are used in a normative sense. Ridge takes this feature to be an improvement over Kratzer's model because it allows for the intelligibility of normative dilemmas, i.e. normative requirements to do the impossible. However, this issue is not particularly relevant in the context of Eklund's challenge, so I am setting it aside here.

¹⁸ Laskowski (2014), Fogal (2016), and Worsnip (2019) have made similar points about the compatibility between contextualism and realism about at least some normative claims.

that guide practical judgments or decisions, and to adopt a normative perspective is to have certain behavioral dispositions with respect to those standards, e.g., dispositions to act in line with certain standards, to issue certain prescriptions, to hold others accountable when they violate those prescriptions, etc.

Putting these elements together, we arrive at a hybrid expressivist view on which authoritatively normative claims express both a normative perspective and a corresponding representational belief. To use the same example, “You ought to give a large part of your income to charity,” if used in an authoritatively normative sense, will express a relational mental state comprising a normative perspective and the belief that donating a large part of your income to charity is highly ranked by any acceptable standard of practical reasoning.¹⁹

On this picture, both authoritatively and non-authoritatively normative concepts can exhibit (CN), but only the former are also referentially normative, partly in virtue of their noncognitive content.

Take non-authoritatively normative concepts first. The semantic analysis of etiquette claims, for instance, will be purely descriptive on Ridge’s view: it will refer to what is highly ranked by certain locally accepted standards of behavior, without involving a normative perspective or any other desire-like attitude directed at those standards.²⁰ But judgments about what is highly ranked by locally accepted standards of behavior also have a certain normative role, at least for the typical concept user: people will standardly rely on such judgments about social practices to guide their actions, criticize others, etc., even if these

¹⁹ Ridge uses the label *metasemantic expressivism* for this view, because it helps explain *why* normative claims have the context-sensitive semantic contents that they do, but Finlay (2014) and Alwood (2016) have argued that we are still dealing here with expressivism as a theory in first-order semantics. Nothing in the present paper hinges on this taxonomical issue.

²⁰ This is compatible with holding that etiquette claims *pragmatically* convey that the speaker has certain desire-like attitudes concerning the relevant standards, at least in typical circumstances. The claim here is that such attitudes are not part of the meaning of etiquette claims, nor do they help explain said meaning.

behavioral dispositions are not constitutively linked to the meaning of etiquette claims. Moreover, this normative role will be roughly the same for all communities using etiquette concepts: that is, etiquette claims will be embedded in similar patterns of behavior and social sanctions in all contexts of use. Thus, etiquette concepts arguably exhibit (CN): if two communities use certain etiquette concepts that have the same normative role, then those concepts are identical.²¹ At the same time, specific standards of etiquette will vary between communities, and therefore the reference of etiquette concepts will vary as well.

For authoritatively normative concepts, in contrast, quasi-naturalism delivers both (CN) and (RN). Again, (CN) is easily secured by expressivism. For example, on this picture, judgments about what is morally wrong have a distinctive role in guiding individual deliberation and interpersonal criticism because they encode normative perspectives, and both we and the community in Alternative use the same concept of wrongness in virtue of this noncognitive content that is constitutively linked with the normative role of “wrong,” even if this content is instantiated by a different normative perspective in each case.

Quasi-naturalism can also vindicate (RN) for authoritatively normative concepts, by treating questions about the reference of terms like *ought* and *wrong* as internal normative questions, to which normative inquiry can establish unique answers that hold true no matter how those concepts are used by some community or other.

A key tool here is deflationism about reference. On a deflationary account, a question such as “What is the reference of ‘wrong,’ for us and in the Alternative scenario?” will be treated as equivalent to the normative question “What kind of actions are wrong, for us and in Alternative?”. To answer this question is to identify the natural features that make actions wrong in the two scenarios, through substantive normative theory.

²¹ N.B.: The claim here is not that these concepts are identical *because* they have the same normative role.

Expressivists have the resources to claim that the same features make actions wrong in all the relevant cases, no matter how any community uses the concept of wrongness. Again, expressivism is not a form of relativism, not even when built into a contextualist framework. Ridge's hybrid expressivism, for instance, provides an account of what it is to *think* that a standard is acceptable, not an account of what makes standards acceptable. On this view, the truth conditions of normative claims do *not* make reference to our normative perspectives, but simply refer to acceptable standards. Moreover, as already mentioned, expressivists can make sense of the idea of objective acceptability—and therefore of objective wrongness, objective goodness, and so on—by identifying a special kind of attitude expressed by objectivity claims, e.g., a normative perspective that condones or rules out certain standards of practical reasoning even with respect to scenarios in which we or others adopted different normative perspectives. By adopting such resilient normative perspectives, we can coherently claim that terms like *wrong* and *good* have a unique objective reference.

To be sure, normative theorizing about what kind of features make actions right, wrong, etc. need not always lead to objectivist conclusions. Indeed, there may be good reasons to adopt relativism about *some* normative issues. I am not claiming, then, that expressivism and deflationism together *entail* Referential Normativity about authoritatively normative concepts, but only that they provide a framework in which first-order normative arguments can be given for the idea that these concepts have a unique objective reference. For instance, once expressivism helps establish that we and the people in Alternative use the same concept of wrongness and are in genuine normative disagreement, utilitarians among us might argue that failing to maximize utility is the fundamental feature that makes actions wrong for both communities, as well as in any other social context. If this argument in

normative theory succeeds, then “wrong” has a unique objective reference: it refers to the property of failing to maximize utility in all the relevant scenarios.²²

Someone might argue that precisely the kind of scenarios imagined by Eklund, in which different communities use authoritatively normative concepts in vastly different ways and seemingly with equal justification, give us reason to reject (RN) with respect to those concepts. In other words, these scenarios still seem to involve an alarming symmetry between the two communities: whatever arguments we might offer for our judgments about the reference of normative concepts, it seems that members of the imagined community will be able to offer similar arguments for their own verdicts on this issue.

I acknowledge that quasi-naturalism still faces a challenge here, but this is a version of the traditional epistemological challenge from disagreement, which all forms of normative realism face, rather than a metasemantic challenge about how to secure (CN) and (RN) for our concepts. Moreover, this epistemological challenge can only arise *after* the metasemantic challenge has been resolved, because it takes for granted that the relevant communities genuinely *disagree* on normative matters, i.e., they use the same concepts to talk about a common subject matter. However pressing this epistemological challenge might be for quasi-naturalists and other realists, this is not Eklund’s challenge anymore: it has nothing to do with alternative normative concepts and the limits of our own conceptual framework.²³

²² To be clear, I do not think it likely that utilitarianism will turn out to be the correct moral theory. This is just an illustration of the idea that normative theorizing can secure a unique objective reference for normative terms.

Another strategy that might allow expressivists to secure an objective reference for at least some normative concepts would be to adopt a constitutivist view on which certain normative truths follow from commitments that are inescapable for any rational agent. On the compatibility between expressivism and constitutivism, see Silverstein (2012) and Ridge (2018).

²³ Street (2011) argues that her Darwinian dilemma for normative realism applies to expressivist quasi-realism as well, and this arguably holds for all epistemological challenges to realism, including the challenge from irresolvable disagreement that I mention here. See Gibbard (2011) and Golub (2017b) for quasi-realist responses

Before moving on to other objections to my proposal, two clarifications are in order. First, the quasi-naturalist defense of (RN) for authoritatively normative concepts does not require Ridge's sophisticated version of hybrid expressivism. Similar arguments for the idea that normative concepts have a unique objective reference can be made using the simple version of hybrid expressivism stated in the previous section, or even in a pure expressivist framework. The reason why I have relied on Ridge's view here is that its contextualist framework allows us to easily distinguish between authoritatively normative concepts, for which both expressivism and (RN) hold true, and relativistic normative concepts, while making good on the plausible idea that terms like *ought* and *good* have core meanings that remain stable across authoritative and non-authoritative contexts of use.

Secondly, I have not claimed that expressivism is the only option for naturalists who want to defend (RN) and rule out the possibility of alternative authoritatively normative concepts. Perhaps there are non-expressivist ways to achieve this. Indeed, the idea that naturalists can treat questions about the reference of normative terms as internal to normative theory has been proposed before in discussions of the Moral Twin Earth problem (see, e.g., Sayre-McCord 1997 or Brink 2001), and some defenders of naturalism even tie the identity of normative concepts to their normative role in a way that resembles the expressivist approach to this issue.²⁴ I am skeptical that any of these naturalist proposals will be able to properly address Eklund's challenge without in effect adopting an expressivist account of

to Street's evolutionary debunking argument and a related epistemological challenge, viz. the demand to explain the reliability of our normative beliefs.

²⁴ Copp (2018), for instance, holds that we should take another community's term "wrong" to mean the same as our "wrong" if they are motivated to act in the same way as us in relation to using that word, even if they otherwise apply "wrong" in very different ways from us. See also Yetter-Chappell and Yetter-Chappell (2013), who suggest that naturalists can adopt a view on which moral concepts are individuated by their action-guiding role and phenomenal character, and therefore the same moral concepts can be used by people who otherwise vastly differ in their dispositions to apply those concepts to things in the world.

normative discourse, but this is a topic for another paper. My goal here has only been to make the positive case that expressivism—and more precisely, a hybrid expressivist view that fits into a contextualist framework—can help vindicate naturalist realism about normativity in the face of Eklund's challenge.

5. The embarrassment of riches problem

In the remainder of this paper, I will address two worries about my proposal. The first is a worry raised by Eklund for anyone who attempts to avoid the problem of alternative normative concepts by defending Referential Normativity: even if normative role determines reference, he argues, there can still be normative concepts with *slightly different* normative roles than our concepts and which refer to something else. This is a problem because the concepts at issue would be intuitively in normative competition with ours, and yet the slight differences in normative role entail that realists cannot rely on Referential Normativity to argue that those concepts must have the same reference as ours. Or so the argument goes.

Expressivist quasi-naturalism is squarely within the target of this objection. As Eklund puts it, discussing Gibbard's version of expressivism:

“One thing Gibbard fails to ask (...) is whether there is, instead of a unique possible concept that is the last ‘ought’ before action, a number of slightly different, non-coextensive concepts each with an equal claim to be the last ‘ought’ before action, and different thinkers can employ different ones among these non-coextensive concepts.” (56-57)

Moreover, Eklund argues, while other realists might respond to this problem by claiming that only some of the competing concepts with slightly different normative roles refer to genuine normative properties—in other words, that normative reality only supplies a reference for

only some of the concepts at issue²⁵—expressivists cannot easily make such a case, given their aversion to metaphysical explanations of facts about truth, reference, etc.

Eklund also mentions another response that realists might adopt in the face of this problem, which seems more congenial to expressivism: “provide a theory of what normative roles are such that expressions associated with different normative roles cannot be in normative competition” (2017, 58). However, he is skeptical that expressivists, or anyone else for that matter, can make good on this idea.

I believe quasi-naturalists can address this worry in a way that involves elements of both of these responses. First, while conceding that concepts with slightly different roles can be in normative competition, quasi-naturalists will insist that such competition is possible only because the normative roles of those concepts largely overlap: that is, there can be genuine disagreement between uses of those concepts precisely because they play a similar role in guiding deliberation and interpersonal criticism and advice, which allows for the possibility of conflicting practical prescriptions.²⁶ For instance, concepts with very different normative roles, such as our concepts of moral rightness and beauty, cannot be in normative competition, at least not in a way that would be relevant to debates about moral or aesthetic realism.

Conversely, to the extent that two concepts do have similar enough roles, their uses will allow for genuine disagreement and mutually intelligible normative debate, even if those concepts are not identical. To take an example, think of our concept of *ought*, and imagine a community that uses a homophonic term with a similar normative role, but with a stronger conceptual connection to motivation and action, akin to our concept of *must*. While these

²⁵ This is what Eklund calls the *normative sparseness* reply.

²⁶ For canonical accounts of disagreement in attitude understood as actual or potential practical conflict, see Stevenson (1944), Ch. 1, and Gibbard (2003), Ch. 4. For a recent articulation of this idea that aims to avoid some problems faced by Stevenson’s and Gibbard’s accounts, see Ridge (2014), Ch. 6.

concepts have slightly different normative roles, their uses can be in genuine disagreement, e.g., if one of us says “We ought to give a large part of our income to charity” and someone in the other community says “We ought to keep most of our income to ourselves”. From an expressivist standpoint, concept identity only matters because it tracks genuine normative disagreement, and such disagreement is possible even in the absence of concept identity.

Finally, such cases of disagreement between uses of concepts with slightly different normative roles pose no special challenge to normative realism, or no challenge having to do with the limits of our conceptual tools, ineffability, or making good on the idea that normative reality privileges some concepts over others. Determining the correct reference of “ought,” both for us and for the other community, now becomes an internal normative question, to be settled with the tools of normative theorizing. That is, we can make a first-order normative case for why our claims involving “ought” are true and the claims of the other community are false, and this is in effect to argue that normative reality privileges our use of “ought”.

7. Is quasi-naturalism a form of ardent realism?

The second challenge I want to briefly discuss is most likely the reason why expressivism has been largely ignored as an option for addressing Eklund's challenge, both in his book and in the ensuing debate. This is the complaint that expressivist quasi-realism, including the naturalist version that I have endorsed in this paper, is not a form of genuine realism.

There are actually two different worries here. One concerns expressivism and its alleged incompatibility with realist commitments. The second worry is about the deflationary framework of quasi-realism. Let me address them separately.

Eklund himself argues that quasi-realism entails a form of relativism, and therefore cannot make good on genuine normative objectivity:

“A quasi-realist view (...) allows one to say that it is true that such-and-such things are bad whenever one has the right attitude, and that when Bad Guy believes something to the contrary, he believes something false. But it is clear that Bad Guy can with equal propriety say the corresponding things about us. Given his attitudes, it is semantically appropriate for him to assert that (say) torturing babies for fun is right, and it is then, on the quasi-realist view, equally appropriate for him to assert that it is true that torturing babies for fun is right; and there are no genuine facts about which he is mistaken.” (2017, 2-3)

Note that this argument concerns the expressivist ingredient of quasi-realism, i.e. the idea that normative claims semantically express desire-like attitudes, rather than its deflationary account of truth, facts, etc.²⁷

This is an instance of a familiar complaint about expressivism, with which I cannot properly engage in this paper. Let me just say, though, that I take this worry to conflate two different questions: whether it is *semantically* appropriate for someone to make a certain normative claim, and whether it is appropriate or correct, *in a substantive normative sense*, for the person in question to make that claim. Or, as Schroeder (2014) puts it, this type of objection to expressivism mistakes the *sincerity* or *assertibility* conditions of normative claims for their *truth* conditions. Expressivism does entail that the sincerity conditions of normative claims depend on the speakers' desire-like mental states, but this is compatible with an objectivist account of normative truths, on which sincere or semantically appropriate normative claims can nevertheless be objectively false. Moreover, quasi-realists can perfectly well claim that there are *genuine facts* about who is correct and who is mistaken in cases of normative disagreement, or even that normative reality favors our ways of valuing and acting,

²⁷ Indeed, in fn. 6, p. 4, Eklund explicitly says that ardent realists *can* be deflationists.

by relying on their deflationary account of facthood and other related notions, on which such seemingly metaphysical claims simply rehearse first-order objectivist commitments.

This is where the second worry about quasi-realism typically arises: deflationism itself, it might be argued, cannot vindicate the existence of genuine normative facts and properties. Copp (2001), for instance, holds that “the central doctrine of moral realism is that moral predicates refer to *robust* moral properties” (38-39, my italics), as opposed to moral properties understood in a merely deflationary sense. Similarly, McPherson (2020) argues that ardent realism requires “an ontologically significant, non-minimalist interpretation of the metaphysical terminology” (11). And these are just two examples of the wide aversion to deflationism among self-described robust realists in metaethics.²⁸

Here is another way to put this worry about deflationism, focusing on the claim that reality itself favors our ways of valuing and acting, which is at the core of Eklund’s challenge. Perhaps quasi-naturalists can make good on this realist idea in some sense, by relying on a deflationary notion of *normative reality* and thereby treating this claim as internal to normative discourse, the argument will go, but they cannot vindicate the metaphysical thesis that reality, in an *external* sense that is independent of any normative perspective, privileges our normative concepts and the ways in which we put them to use. And it is this latter thesis that defines ardent realism.²⁹

²⁸ One option available to quasi-naturalists might be to adopt a non-deflationary account of some of the relevant metaphysical and semantic notions, such as truth or property. See, e.g., Ridge (2014), who argues that his version of quasi-realism does not require deflationism and can be combined with a correspondence conception of truth, albeit one on which truth claims are understood as normative and given an expressivist account. I am skeptical, though, that relying on this conception of truth would allay the worries of those who think that quasi-naturalism cannot address Eklund’s challenge because it is not a form of genuine realism. Their complaints are likely to be redirected at the expressivist account of truth itself.

²⁹ Thanks to an anonymous reviewer for suggesting this way of framing the worry about the deflationary framework of quasi-naturalism.

I cannot do justice to the debate about deflationism and the question of realism in the present paper either, but let me make two brief points.

First, here is how things look from a deflationary standpoint: normative truths and facts understood in a deflationary sense are as robust as they can intelligibly be, and the same goes for the notion of normative reality.³⁰ So deflationism is compatible with full-blown realism about any domain of facts, at least by its own lights. The fact that there is a dispute in metametaphysics about whether deflationism is the right account of truth, facts, reality, etc. does not entail that any metaethical view that adopts deflationism is thereby anti-realist, and more generally does not affect the question of realism in any domain of discourse.³¹

Secondly, this deflationary approach to metaphysical issues does entail that there is no sense to be made of the idea that reality favors our ways of valuing and acting *except* as a claim internal to normative discourse, equivalent to the claim that everyone ought to use normative concepts as we do or to some similar normative claim. But this is as it should be, quasi-naturalists will insist: we cannot intelligibly demand that reality understood in a *non-normative* sense favor our normative outlook, and it is a virtue of quasi-naturalism that it avoids such confused ideas by interpreting all metaphysical claims about normative facts, reality, or objectivity as organic parts of normative theory.³²

³⁰ Of course, people can coherently reject deflationism and propose a different interpretation of metaphysical claims about normative truths, facts, reality, etc. But if deflationism *is* indeed the correct account of the meaning of terms like “true,” “fact,” etc., then there is in fact no intelligible way of using such words in a more robust metaphysical sense. Thanks to Max Hayward for pressing me on this issue.

³¹ For example, few people would dispute that someone can be a deflationist about truth, facts, etc. and yet a scientific realist. It is not clear to me why deflationism is treated with much more suspicion by normative realists compared to realists in other domains.

³² As Ridge (2018, 2960) puts it: “The idea that we might somehow stand outside our normative commitments and worry about whether they correspond to some putative normative reality which we can conceptualize independently of making any specific normative judgments is entirely alien to the quasi-realist approach.” And this idea that realist-sounding claims about objectivity or mind-independence can only be understood as internal

To be sure, there are significant differences between quasi-realism and other forms of normative realism: most importantly, an explanatory contrast that bears on central issues in metaethics. Quasi-realists explain the content and the distinctive nature of normative concepts in terms of their non-representational, practical function, rather than by appealing to representational relations between such concepts and normative entities. Indeed, this explanatory contrast is the reason why many quasi-realists themselves see their view as distinct from realism (e.g., Blackburn 1993, Gibbard 2003).³³

However, as I have argued in previous work (Golub 2017a, 2021), there seems to be no explanatory conflict here between quasi-realism and anything recognizable as a general notion of normative realism, at least not if we adopt a deflationary standpoint. For instance, we cannot define realism in terms of explanations of normative meaning that ascribe a substantive role to truth and reference relations understood as general semantic notions, because such explanations are incompatible with deflationism about truth and reference. In other words, if we take deflationism seriously, we will not see these explanations of meaning as indicating some robust metaphysical commitments (“robust realism”) that quasi-realists should reject, but rather as the product of confusion about the nature of truth and reference.

For these reasons, I believe quasi-naturalists should think of their view as a form of realism pure and simple, and the same goes for anyone who accepts that the quasi-realist project *might* succeed. And if my arguments in this paper are correct, quasi-naturalism might offer the best way to vindicate ardent realism in the face of the problem of alternative

to normative discourse is a recurrent theme of Blackburn’s work. See, e.g., Blackburn (1993: 156-157, 173), among many other places.

³³ For different versions of this idea that quasi-realism and genuine realism diverge in their explanatory commitments, see also Dreier (2004) and Simpson (2018), among others.

normative concepts precisely because it rejects a representationalist account of normative meaning and normative concepts.³⁴

References

- Alwood, A. (2016). Should Expressivism Be a Theory at the Level of Metasemantics? *Thought: A Journal of Philosophy*, 5, 13–22.
- Baker, D. (2017). The Varieties of Normativity. In T. McPherson and D. Plunkett (eds.), *The Routledge Handbook of Metaethics*, Routledge, 567-581.
- Bex-Priestley, G. (forthcoming). Expressivists Should Be Reductive Naturalists. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 17. Oxford University Press.
- Blackburn, S. (1984). *Spreading the Word*. Clarendon Press.
- . 1993. *Essays in Quasi-Realism*. Oxford University Press.
- Boyd, R. (1988). How to Be a Moral Realist. In G. Sayre-McCord (ed.), *Essays on Moral Realism*, 181–228. Cornell University Press.
- Brink, D. O. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge University Press.
- Brink, D. O. (2001). Realism, Naturalism, and Moral Semantics. *Social Philosophy and Policy*, 18.
- Chrisman, M. (2016). *The Meaning of ‘Ought’: Beyond Descriptivism and Expressivism in Metaethics*. Oxford University Press.

³⁴ Many thanks to Graham Bex-Priestley, Darren Bradley, Daniel Elstein, Will Gamester, Max Hayward, Pekka Väyrynen, Richard Rowland, Jack Woods, two anonymous reviewers, and an audience at the 2021 Value in Language Workshop for comments and discussion on previous versions of this paper. Research for this paper was funded through the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 837036.

- Copp, D. (2001). Realist-Expressivism: A Neglected Option for Moral Realism. *Social Philosophy and Policy*, 18, 1–43.
- Copp, D. (2005). Moral Naturalism and Three Grades of Normativity. In P. Schaber (ed.), *Normativity and Naturalism*, Ontos-Verlag, 7-45.
- Copp, D. (2018). Realist-Expressivism and the Fundamental Role of Normative Belief. *Philosophical Studies*, 175, 1333–1356.
- Dowell, Janice J. L. (2016). The Metaethical Insignificance of Moral Twin Earth. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 11. Oxford University Press.
- Dreier, J. (2004). Meta-ethics and the Problem of Creeping Minimalism. *Philosophical Perspectives*, 18(1), 23-44.
- Dunaway, B. & McPherson, T. (2016). Reference Magnetism as a Solution to the Moral Twin Earth Problem. *Ergo: An Open Access Journal of Philosophy*, 3.
- Dworkin, R. (1996). Objectivity and Truth: You'd Better Believe It. *Philosophy and Public Affairs*, 25, 87–139.
- Eklund, M. (2017). *Choosing Normative Concepts*. Oxford University Press.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford University Press.
- Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Fogal, D. (2016). Confusion of Tongues: A Theory of Normative Language, by Stephen Finlay. *Ethics*, 127, 281–288.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press.
- Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press.

- Gibbard, A. (2011). How Much Realism? Evolved Thinkers and Normative Concepts. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 6, 33-51.
- Golub, C. (2017a). Expressivism and Realist Explanations. *Philosophical Studies*, 174, 1385–1409.
- Golub, C. (2017b). Expressivism and the Reliability Challenge. *Ethical Theory and Moral Practice*, 20(4), 797-811.
- Golub, C. (2021). Representation, Deflationism, and the Question of Realism. *Ergo*, 7(36).
- Hare, R. M. (1952)
- Horgan, T. & Timmons, M. (1991). New Wave Moral Realism Meets Moral Twin Earth. *Journal of Philosophical Research*, 16, 447–465.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford University Press.
- Kratzer, A. (2012). *Modals and Conditionals: New and Revised Perspectives*. Oxford University Press.
- Laskowski, N. (2014). How to Pull a Metaphysical Rabbit Out of an End-Relational Semantic Hat. *Res Philosophica*, 91, 589–607.
- Lenman, J. (2007). Disciplined Syntacticism and Moral Expressivism. *Philosophy and Phenomenological Research*, 66(1), 32-57.
- Majors, B. (2006). Quasi-Naturalism and Moral Reality. *Ratio*, 19(1), 64–76.
- McPherson, T. (2011). Against Quietist Normative Realism. *Philosophical Studies*, 154(2), 223-240.
- McPherson, T. (2020). Ardent Realism Without Referential Normativity. *Inquiry*, 63, 489–508.
- Moore, G. E. (1903). *Principia Ethica*. Dover Publications.
- Plunkett, D. (2020) Normative Roles, Conceptual Variance, and Ardent Realism

About Normativity. *Inquiry*, 63, 509–534.

Ridge, M. (2007) Ecumenical Expressivism: The Best of Both Worlds? In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 2, 51–76.

Ridge, M. (2014). *Impassioned Belief*. Oxford University Press.

Ridge, M. (2018). Meeting Constitutivists Halfway. *Philosophical Studies*, 175(12), 2951–2968.

Sayre-McCord, G. (1997) ‘Good’ on Twin Earth. *Philosophical Issues*, 8, 267–292.

Scanlon, T. M. (2014). *Being Realistic About Reasons*. Oxford University Press.

Schroeder, M. (2008). *Being For: Evaluating the Semantic Program of Expressivism*. Oxford University Press.

Schroeder, M. (2013). Tempered Expressivism. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 8. Oxford University Press.

Schroeder, M. (2014). Does Expressivism Have Subjectivist Consequences? *Philosophical Perspectives*, 28, 278–290.

Silverstein, M. (2012). Inescapability and Normativity. *Journal of Ethics and Social Philosophy*, 6(3), 1–27.

Simpson, M. (2018). Solving the Problem of Creeping Minimalism. *Canadian Journal of Philosophy*, 48(3–4), 510–531.

Street, S. (2011). Mind-Independence Without the Mystery: Why Quasi-Realists Can’t Have it Both Ways. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 6, Oxford University Press, 1–32.

Toppinen, T. (2013). Believing in Expressivism. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 8, Oxford University Press.

Väyrynen, P. (2018) A Simple Escape From Moral Twin Earth. *Thought*, 7, 109–118.

Worsnip, A. (2019). 'Ought'-Contextualism Beyond the Parochial. *Philosophical Studies*, 176, 3099–3119.

Yetter-Chappell, H. & Yetter-Chappell, R. (2013) Mind-Body Meets Metaethics: A Moral Concept Strategy. *Philosophical Studies*, 165(3), 865–878.