

Non-Cognitivism and Roads Untraveled

Matt Bedke | draft for CHillMeta Workshop 2016

1 Introduction

In this paper I want to re-examine the options for a non-cognitivist metaethic. As I understand it, non-cognitivism has two core commitments.

Non-descriptivism: Normative components to thought/language do not have representational or descriptive content.

Action-guiding: There is something intrinsically action-guiding or practical about normative thought/language that we would like meta-normative theory to explain.

A full defense of non-cognitivism needs to 1) defend these core commitments, 2) provide a theory of normative thought/language that satisfies them, and 3) adequately address the key explananda that drive metaethical theory, such as the possibility of normative disagreement and error, the bundle of logical and semantic properties highlighted by Frege-Geach problems, the supervenience of the normative on the non-normative, etc. I will not take up the first task here. Instead, I want to craft novel versions of non-cognitivism that satisfy its core commitments, and that promise to adequately address the other explananda.

Of course, the most promising version of non-cognitivism to date is expressivism. But the view has enough liabilities to motivate new approaches. According to one popular conception of the view, expressivism is a two-part view about semantic values. Part one: the semantic value of an atomic sentence is the state of mind the sentence is used to express. Part two: unlike many descriptively contentful sentences, the (or one) state of mind expressed by atomic normative sentences is conative – something other than a prosaic belief.¹ This conception of expressivism locates the view in a fairly standard semantic framework, one that posits characters, contexts, contents (/semantic values), circumstances of evaluation and

¹ See, e.g., Rosen (1998, p. 387) and Schroeder (2008a, p. 33).

extensions to model the compositionality and truth-functionality of language. One huge problem, however, is that modeling compositionality and truth-functionality is extremely difficult when semantic values are motely. Hence the well-known Frege-Geach problems for expressivism.

More recently it has been suggested that expressivism should be understood as a meta-semantic hypothesis. On that view sentences get their semantic values (whatever they are) in virtue of expressing states of mind.² The hope is that we can go ahead and assign semantic values as needed to model compositionality—we are not constrained to assign a mental state for atomic values, for example—so long as the meta-semantics of attitude expression can explain *how* those bits of language get the semantic values they get. Though sympathetic with this shift in perspective, I have my reservations. For one, it is not clear to me that we have a grip on some notion of attitude expression capable of bearing explanatory weight in a meta-semantic theory. Second, even if there is some notion of attitude expression that can bear that weight, the meta-semantic move does not directly address non-cognitivism's core commitments. One key question is whether normative language has representational content—whether it purports to represent facts and properties—and it is not at all clear that meta-semantic expressivism speaks to the issue. The obvious ways of supplementing meta-semantic expressivism so that it avoids representational content (e.g., assign conative mental states as semantic values) face the same difficulties as expressivism *qua* hypothesis about semantic values.

Perhaps, then, expressivism is part and parcel of a more radical pragmatic theory of meaning, sometimes known as global expressivism or global anti-representationalism.³ The main problem with that approach is that it threatens to smooth over important distinctions that traditional non-cognitivists wanted to keep front-and-center. For example, they want to

² See, e.g., Chrisman (2012, pp. 323-30), and Ridge (2014, pp. 102-11, 124-31).

³ See, e.g., Price et al. (2013).

say that, *unlike* many non-normative declarative sentences, normative declarative sentences do not represent, at least not by virtue of their normative vocabulary.⁴ It is just not clear how to draw the representational, non-representational lines in the right places within a global pragmatic theory. Moreover, to the extent that such theories eschew formal semantic frameworks, they risk losing some of their explanatory power.

That is a bit of a romp through some of my concerns with expressivism. In any event, it would be nice to have alternative versions of non-cognitivism without taking a step back to emotivism or prescriptivism. Here is the plan. Sections 2 and 3 draw from theories of vagueness. There, the idea is that non-cognitivist can satisfy the non-descriptive hypothesis with an admittedly odd idea: Thin normative predicates are *maximally vague*. That is, all cases that are eligible for falling under these predicates are borderline cases, at least when speaking of descriptive dimensions to meaning.

Sections 4 and 5 draw on theories of use-conditional semantics. There, the idea is that non-cognitivist can satisfy the action-guiding aspects to meaning and clear certain explanatory hurdles *without* positing descriptive content. The basic idea is to assign a special sort of (Kaplanian) character—or cognitive role—to the favored normative expressions, one that generates speaker-variant, and conative-attitude-dependent, contents for those expressions. This character introduces and exploits a distinction between two different relations language can bear to its content – a descriptive relation, and a non-descriptive relation. The distinction allows non-cognitivists to avoid a *descriptive relation to contents* (*/semantic values*) without having to introduce any *novel type of content* (e.g., mental states, hyperdecided states, etc.). This important distinction can then be replicated when it comes to

⁴ To be sure, pragmatic approaches can marshal some resources to distinguish intuitively representational language from intuitively nonrepresentational language. But it is difficult to draw the lines in the right places. For example, Huw Price distinguishes language that has the function of tracking facts from language that does not, but this leaves no space for the possibility of fragments of language that *purport* to represent even though they lack the purpose tracking facts (as error theorists might say of normative discourse).

norms of assertion. It can be semantically appropriate to assert a normative sentence only when certain worldly conditions obtain, but that does not entail that the assertion is in the business of representing those conditions as obtaining.

That is a bit quick, and it will take some time to unpack the main ideas. In the end we will develop three interrelated non-cognitivist theories: 1) combine maximal descriptive vagueness with a *pragmatic* theory of how the favored expressions are action-guiding, 2) combine maximal descriptive vagueness with a *semantic* theory of how the favored expressions are action-guiding, one that assigns non-descriptive (perhaps better, *non-described*) contents to the favored expressions, and 3) drop maximal descriptive vagueness and make do with non-descriptive contents.

Of course, I cannot hope to address all of the explanatory hurdles facing non-cognitivism. So I focus on what are arguably the most imposing ones, and certainly the ones that drive much of the recent literature – the bundle of semantic, logical and inferential problems that often fly under the banner of Frege-Geach problems.

2 Why Vagueness Should Be Interesting to Non-Cognitivists

Let me begin with a few remarks about vagueness so that we can see why non-cognitivists might be especially interested in the phenomenon. Set normativity aside for a moment and consider a vague predicate like ‘is a heap’ and its application to a spectrum of different quantities and arrangements of grains of sand. The spectrum begins with a single grain of sand, and each successive case adds a single grain until one ends the spectrum with a pile as high as a person. It is clear that cases at the end of the spectrum are heaps. It is just as clear that cases at the beginning of the spectrum are not heaps. But somewhere between the extremes there are hard cases, the so-called *borderline cases*. Though it is controversial what to say about such cases, I take it that the meaning of a natural language predicate like ‘is a

heap' is not so precise so as to fix a point in the spectrum where everything to one side falls under the predicate and everything to the other side does not.

There are lots of things one might say within formal semantics to make sense of borderline cases. But less formally, it looks like the sentence “*a* is a heap” does describe *a* as being some way – it describes *a* as being *unlike* those piles of sand that are clearly (/determinately) not heaps. Also, “*a* is not a heap” describes *a* as being some way – it describes *a* as being *unlike* those piles of sand that are clearly (/determinately) heaps. But neither “*a* is a heap” nor “*a* is not a heap” seem to describe *a* as being similar to or different from the borderline cases. In fact, they seem to leave open whether *a* is like a borderline case or not. Here is one way to put the thought: Neither sentence *rules out* the possibility that *a* is a borderline case.

Insofar as failure to rule out borderline cases is a kind of failure to describe what things are like—a failure to discriminate among possibilities—this aspect of vagueness should be very interesting to non-cognitivists. For one way for a term to lack descriptive content would be for it to be *maximally vague*. A maximally vague predicate would just have borderline cases, and sentences predicating it would thereby fail to rule out *any* possibilities, at least in terms of its conventional content.⁵

⁵ Though not an obvious way to go with semantic theory, the approach has some pedigree. For example, Joshua Gert (2007, pp. 92-95; 2012, pp. 51-6) suggests that ‘is funny’ is “maximally vague”. But by this he means that no application of ‘is funny’ is determinately true. He does allow for some applications to be determinately false. This isn’t the radical maximal vagueness I advocate for thin normative terms. And his way of fixing extensions, and fixing a domain of attitude expression, differs from my approach here. He takes shared responses to normative cases to fix speaker-invariant content, and divergence in responses to define a borderline region where expressivism might be apt. This makes the degree to which the discourse is cognitive to be the degree to which it is not expressivist. As we shall see, the views developed here are very different, though Gert and I are motivated by similar concerns, and we both think vagueness can help.

Lenhart Åqvist (1964) presents the view that all value properties are “totally vague”, and so all propositions asserting goodness or badness are neither true nor false. His argument turns on Moorean doubts about analyzing value naturalistically. He also notes that the theory

A further reason to be particularly interested in vagueness is that vague discourse, even when it is applied to borderline cases, completely avoids Frege-Geach problems. Consider again borderline cases for non-normative predicates. Take a borderline case of a heap of sand and name it Heapish. Take another borderline case of a heap just like Heapish except that it has one extra grain of sand piled right on top. Name it Heapy. Now consider the following argument.

- 1v. If Heapish is a heap, then Heapy is a heap.
- 2v. Heapish is a heap.
- Cv. Heapy is a heap.

Let us begin with equivocation and validity. The meaning of ‘heap’, though vague, does not equivocate across the premises. It is a meaning we can spell out in terms of contents at contexts of use, and in this case we are assuming the content does rule out some possible ways Heapish might be, and the content does not vary across premises. Further, this is clearly a valid argument. However we assign truth-values to vague sentences, we can say that *if* the premises are true then the conclusion must be true. The leading ways of composing truth-values for vague language agree on this.⁶ I should add that validity under supervaluationism

would need to be supplemented to satisfy the practical, action-guiding aspects of the language, but does not offer that supplementation (as I do here).

Last, Stephen Schiffer (2002) presents the view that all moral propositions are neither determinately true nor determinately false. His reasons for holding the view seem to turn on i) the possibility of rationally irresolvable disagreements on any moral issue (which is something that weighs with Gert, too), ii) the inexplicability of how one could have *a priori* knowledge of moral principles, and iii) his contention that, for any moral principle, one could “vague believe” it to degree .5. He also does not supplement his view to account for practical, action-guiding aspects of the discourse.

Thanks to Giulia Pravato, who brought some of this material to my attention after I wrote up my own ideas. She is working on similar ideas and defends a fascinating theory of the “irresolvable vagueness” of normative terms.

⁶ This is only one argument form, of course, but armed with Łukasiewicz or Kleene truth tables one can see how the points generalize to other argument forms. Supervaluationism has the benefit of preserving all and only classical validities so long as one does not include a ‘determinately’ operator in the language. A reader has worried that the absence of valid formulae on the strong Kleene tables entails that arguments with vacuous premises are

would be the preservation of supertruth, and it is clear that, for every admissible precisification of ‘heap’ under which the premises are true, the conclusion is also true.^{7,8}

As for the inferential moves of reasoning subjects, if one accepts (1v) and (2v) but denies (Cv)—i.e., one accepts ($\sim Cv$)—one thereby makes a logical mistake. For it cannot be that (1v), (2v) and ($\sim Cv$) are all true, and to accept a sentence is to think it true. Further, if one is *justified* in accepting (1v) and (2v), and has no greater justification for denying (Cv), one has justification for accepting (Cv). There is no magic here. This is how justification transmits in valid arguments. The only problem here is if there is some bar to justifiably accepting an indeterminately true sentence. I see no such bar. Often the fact that some item is a borderline case is not transparent to us, and there can be good reasons to think some grains of sand amount to a heap even when they constitute a borderline case. So we seem to get standard inferential roles wherever justified acceptance is a live possibility.

In short, theories of vagueness have the resources to completely avoid Frege-Geach problems *even when we restrict our attention to the borderline cases* of Heapish and Heapy.

This in mind, let us turn to the following argument, where we assume ‘wrong’ is a thin normative term, and we postulate that thin normative terms are maximally vague.⁹

1. If tormenting the cat is wrong, then tormenting the dog is wrong.
2. Tormenting the cat is wrong.

invalid. That is not right. Validity here is truth of the conclusion for all Kleene interpretations under which all the premises are true. A *modus ponens* argument is still valid on this view. So even though $(p \ \& \ (p \supset q) \supset q)$ is not a valid formula, (1) p ; (2) $p \supset q$; (C) q , is a valid argument. The only Kleene assignment under which the premises are true is one where the conclusion is true. It matters not whether the conditional is vacuously true.

⁷ For the classic statement of supervaluationism, see Fine (1975).

⁸ See Cobreros (2011) for discussion and suggestions for deductive systems with the ‘determinately’ or ‘definitely’ operator.

⁹ For the initial Frege-Geach problems, see Geach (1960, 1965) and Searle (1969, pp. 136-141). For some nice discussions see Dreier (1996), van Roojen (1996) and Schroeder (2008b). See also Baker and Woods (2015), Blackburn (1984, ch. 6, 1998, ch. 3), Charlow (2014), Gibbard (2003, chs. 3-4), Horgan and Timmons (2006, 2009), Schroeder (2008a), and Schwartz & Hom (forthcoming).

C. Tormenting the dog is wrong.

We can say the same thing for this argument as we said for the one above. The main difference here concerns the contents of the sentences. For application of a non-maximally vague predicate, the content will rule out some possibilities. Here, however, the hypothesis is that *no* possibilities are ruled out. There are various ways of modeling this in the formal semantics. One option is to say that these sentences each express the set of all possible worlds, leaving nothing out. Or, alternatively, they express sets of sets of possible worlds, but again where no possible world is left out. The basic idea is that normative sentences will not be discriminating at all when it comes to representing how things might be.

Is there any reason to think that maximally vague predicates would fair worse than standard vague predicates applied to borderline cases when it comes to Frege-Geach problems? No. To see this, ask what would happen to the argument from (1v) and (2v) to (Cv) if we were to make ‘heap’ more and more vague; i.e., what happens if we expand its borderline region? Let us imagine taking the cases in the extension and anti-extension of the term as-is, and one-by-one placing them in the borderline region. (We can imagine the descriptive-content-fixing rules, or whatever it is that helps to fix content, slowly shifting.) At any given point in the process does there result a term that lacks embeddability into valid arguments? No. Put the point this way: There is no reason to think that logical inconsistency in accepting (1v) and (2v) while denying (Cv) turns on the size of the borderline region.

When it comes to predicating of a borderline case we can fill in the story in a number of ways. For example, when someone says “Heapish is a heap” or “Tormenting the cat is wrong”, we *could* say that she makes a *semantic mistake* by applying a term to an item not within the extension of the term. But I think it is more plausible to say that she does not make a semantic mistake, and neither does a person who says “Heapish is not a heap” or “Tormenting the cat is not wrong”. Borderline cases are cases where the semantic rules leave

predication optional – where predication of ‘is F’ and ‘is not F’ are both permissible and so both semantically tolerable. Let us call the conditions that are neither sufficient for falling under a predicate (application being semantically required) nor insufficient for falling under the predicate (application being semantically forbidden) *merely adequate conditions*.

This is not to rule out *disagreement* over whether Heapish is a heap or not or whether tormenting the cat is wrong or not. It is simply to deny that these speakers are making a certain kind of mistake, *so any disagreement among them cannot be cast as a violation of the semantic rules*. Compare: Someone can use a racial epithet without violating the semantic rules concerning how that term is to be used, we can acknowledge the semantic propriety of his use, and yet we can still take issue with what he says and in some intelligible sense disagree with him. For normative terms, the idea is that application of normative terms is semantically permissible and merely permissible across the term’s entire domain of application.¹⁰ Again, this is not meant to rule out disagreement *tout court*. I take this up in the last section.

Though maximal vagueness is an unfamiliar way of thinking of how a predicate can lack descriptive content, it does have benefits. Unlike expressivism, we do not assign a motley assortment of mental states as semantic values, and so we do not rely on an ideationalist semantics, nor do we introduce composition of attitudes. We can get non-descriptivism without the Frege-Geach problems that beset expressivism. At least, that is the hope.

3 Prima Facie Problems

¹⁰ See Shapiro (2006, pp. 10-12) and those he cites. Sainsbury raises a similar issue with the “mandatoriness” of predication (1990, p. 259). When speaking of the optionality of applying ‘F’ to *a* it should be understood that we are speaking of predication made in isolation. When coordinating predications with a single predicate we need something like Kit Fine’s penumbral connections (1975, p. 271). See note 22.

The obvious problem with the view so far is that there seem to be clear, non-borderline cases of thin normative statuses. How about torturing children for fun? That does not seem to be a borderline case. It seems wrong. It seems true that it is wrong.¹¹

Relatedly, would not maximal vagueness stand in the way of justifiably accepting any normative sentence, and so of inferential role? Not only does it seem true that torturing children for fun is wrong, but I think I have good reason to believe it. How does that square with being a borderline case and there being no sufficient descriptive conditions for falling under ‘wrong’? In the case of ‘bald’, we can see how the status of being borderline might not be cognitively transparent, and how one might have good evidence that a borderline bald person is clearly bald. But if a thin normative term like ‘is wrong’ is *maximally* vague it is harder to see how it comes to pass that we take any considerations to be good evidence of wrong-doing.

These concerns are somewhat mitigated if predicating of a borderline case is semantically permissible. At the very least we can say that it is no semantic mistake to assert that torturing children for fun is wrong. Then again, asserting that torturing children for fun is *not* wrong is not a semantic mistake either. So there remains the question of why we accept the normative sentences we in fact accept, and how we could justifiably do so if they are all about borderline cases.

Last, though we can see how maximal vagueness satisfies Non-descriptivism, it does not yet speak to Action-guiding. Perhaps the simplest way to satisfy Action-guiding is to locate the action-guiding qualities of normativity in the *pragmatics of language use*. On that view, when a speaker uses normative language her conative attitudes would be calculable based in part on conversational norms and so not solely on the basis of the conventional content of what is said. That yields what I will call *VAP Non-cognitivism*, or vagueness and

¹¹ Roy Sorenson criticizes Åqvist’s theory on the grounds that some cases seem like clear cases, e.g., Hitler was not a borderline case of an evil man (1990, p. 11).

pragmatics non-cognitivism. I will not develop this version of non-cognitivism much farther, for I think there are serious concerns however it gets developed. For one, it does not help to solve the problems of maximal vagueness just noted. We think there are non-borderline cases of being wrong, and we think we are sometimes justified in thinking some act wrong. The pragmatics of attitude expression will not help to deliver those results. Second, it is hard to see how use of maximally vague terms would pragmatically communicate one's conative attitudes.¹² Third, there are general problems with locating the action-guiding qualities of normativity in pragmatics. The semantic-pragmatic divide is tricky, but in broad strokes if someone sincerely says, "Tormenting the cat is wrong" I do not figure out that s/he has a negative attitude toward tormenting the cat (or toward some of that act's properties) by appeal to rules of good conversation and communication. Instead, the communication of negative attitude seems encoded in the conventional meaning of the term such that one who says, "Tormenting the cat is wrong, but I have nothing against it" is making a semantic mistake. She is using the wrong term to render a judgment about tormenting the cat if she lacks *any* relevant negative attitude. She is violating the semantic rules of correct usage. In this respect the case is similar to those cases where one uses a pejorative term for a racial minority, but lacks any negative attitudes toward those minorities. It is not just that pragmatically he conveys a negative attitude he does not have. He is abusing the word. Relatedly, and fourth, it is widely thought that normative *thoughts* have some special motivational role that standard, non-normative beliefs lack. We might attribute mental states to beings that speak an untranslatable language, and it is likely that we would take practical, motivational roles seriously when attributing normative mental states. If so, the pragmatics of language use is the wrong toolbox to reach for.

¹² Giulia Pravato is working on one theory of how this might happen. Her view relies on the idea that thoughts employing vague concepts involve accepting precisifications of those concepts. It is the acceptance of the precisification that is then communicated with language employing vague terms.

Thankfully, VAP Non-cognitivism is not the only way to go. If we draw on use-conditional semantics we can address some of the problems of maximal vagueness and satisfy Action-guiding in a more satisfactory way.

4 The Semantics of Guiding Action

Recall that one expressivist approach starts by assigning mental states to atomic sentences in assertoric contexts – the mental states that are “expressed”. To get the hallmarks of non-cognitivism out of this setup one assigns non-cognitive attitudes to certain atomic normative sentences. “Tormenting the cat is wrong” would presumably get assigned a negative non-cognitive attitude toward Tormenting the cat (or toward some of its properties). The sentence thereby *expresses* that non-cognitive attitude and does not *describe* Tormenting the cat in any way.

The approach effectively rolls the dice on a controversial, psychologized theory of meaning for *all* sentences just to open up *one way* in which a sentence might express non-cognitive attitude. That strikes me as overkill. But the move was a well-intended improvement over speaker subjectivism. According to speaker subjectivism, normative sentences *report on* the speaker’s attitudes, or, in a more sophisticated variety, they report on certain properties picked out by the speaker’s attitudes.¹³

Either way, under speaker subjectivism normative sentences are in the business of describing how things are, and different speakers will presumably *describe different things* depending on their attitudes. This does a nice job of predicting and explaining certain action-guiding qualities. But it does not satisfy Non-descriptivism, and the actual theory of descriptive content generates a host of familiar problems (e.g., how to capture normative disagreement, how to recover certain counterfactuals like “If I were all for bear baiting, it would still be wrong”). The expressivist thinks the key lies in a distinction between *reporting*

¹³ Cf. Dreier (1990, pp. 19-20; 1999, pp. 564-70).

one's attitudes and *expressing* them. But they get attitude expression by modifying the classical theory of semantic content. The leading modifications are to either make the contents of sentences—all sentences—mental states, or to enrich content to include things like hyperdecided states (in addition to classical contents, like sets of possible worlds).

4.1 Contents: Descriptive vs. Non-Descriptive

I want to suggest that expressivism got off on the wrong foot there. There is another way to modify speaker subjectivism. We could leave the contents assigned by speaker subjectivism alone, thereby securing some interesting action-guiding qualities and retaining conventional views of semantic contents. We need only reject a key assumption: If a sentence has a content, assertoric uses of that sentence are in the business of describing that content, or the language is reporting that things are as the content represents, or some such. The alternative is to say that, for some language with content, the language is not in the business of describing that content. Put another way, assertoric uses of those sentences are not attempts to put their content at issue as accurately representing how things are.

There is an elegant way of drawing the desired contrast here. If semantically correct language use can be captured by semantic rules that we grasp (more or less) when we come to understand the meaning of the language (its character or cognitive role), we can capture a distinction between language that describes its content and language that does not in the rules of use. Let me suggest that content that is described, or *descriptive content*,¹⁴ is fixed by what we can call *for use rules*, or f-rules. To illustrate, we can formulate the speaker subjectivist use rule for 'is wrong' like this:

¹⁴ It is important that descriptive content is not a special kind of content. It is not a hyperdecided state as opposed to a possible world. No, it is prosaic content. It is only non-descriptive insofar as it is the content of a bit of language that bears a non-descriptive relation to that content.

Predicate ‘is wrong’ of an action, *A*, for exhibiting dthat(a property the speaker *all in* disapproves of in action).¹⁵

‘For’ is a key term. It signals that there is something about the rules for using the term that makes the predicate *stand for something*. It imbues the predicate with not just content, but *descriptive* content.

This in place, we can see a rather minimal modification to semantic theory to avoid speaker subjectivism while retaining a central role for conative attitude. We just need to modify the use rule so it is not an f-rule. A nice contrast is to posit a use rule that says *when* to use a predicate, but without saying that the predicate is used *for* anything (i.e., not as a sign of anything, or as standing for something). Let me suggest this for ‘is wrong’:

Predicate ‘is wrong’ of an action, *A*, when *A* exhibits dthat(a property the speaker *all in* disapproves of in action).¹⁶

This is subtly but importantly different from the subjectivist’s use rule. Here, we use ‘is wrong’ *when* certain conditions obtain, conditions picked out by our attitudes, but not *for* those conditions. This allows us to retain the content suggested by speaker subjectivism while dropping the reporting function of the language.¹⁷ We simply think of the content of a predicate in terms of its *use conditions*¹⁸—those conditions that have to be exhibited for one to use the term in semantically appropriate fashion—and then distinguish use conditions that the language *describes* or *represents* as being the case from use conditions that are not

¹⁵ I make use of Kaplan’s ‘dthat’ operator, which we can think of as a way of encoding the character of the predicate ‘is wrong’. This allows us to make use of speaker attitude in fixing content without necessarily placing speaker attitude in the content, as sophisticated versions of speaker subjectivism would like to do.

¹⁶ Similarly, consider this rule for ‘practical reason against’:
Predicate ‘is a practical reason against’ of a consideration, *R*, concerning an action, *A*, when *R* is dthat(a consideration the speaker *pro tanto* disfavors in actions like *A*).

¹⁷ Smith and Stoljar (2003) make similar distinction when they respond to Jackson and Pettit’s (1998) contention that conventional agreements to use terms *when* certain conditions obtain are agreements to *report* on the presence of those conditions.

¹⁸ For a nice collection of use-conditional essays, see Gutzmann and Gärtner (2013).

represented or described as being the case. Following this lead, we get a bifurcation of key semantic terms. *F-rules* will be associated with *descriptive* content, sufficient *descriptive* conditions for falling under predicates, etc., whereas *when* use rules—or *w-rules*—will be associated with non-descriptive content, sufficient non-descriptive conditions, and so forth.¹⁹

It should not be too controversial that there are terms whose semantically appropriate use turns on the presence of certain conditions, but where those terms are not used to describe or represent those conditions. ‘Ouch’, ‘bah humbug’, ‘wow’, and ‘damn’ are good examples. If I say ‘ouch!’, you can bet that I am in pain, but I am not thereby describing myself as being in pain. You can accuse me of misusing the term if I am not actually in pain, but you cannot accuse me of misdescribing things. Further, the misuse looks like a violation of semantically encoded meaning, not a pragmatic *faux pas*.

The terms just discussed probably have no descriptive content. Other terms likely have both descriptive content and non-descriptive content. Consider the distinction between formal and informal second personal pronouns in many languages. In Spanish, ‘tú’ is the informal ‘you’ while ‘usted’ is the formal ‘you’. If I say “Tú me debes diez dólares” (you owe me ten dollars) this indicates but does not describe that we have a familiar and informal relationship, a non-descriptive condition, and it also contributes you as part of the described content. If I say “Usted me debe diez dólares” (you owe me ten dollars), this indicates but does not describe a more formal relationship, and also contributes you to descriptive content. Focusing on the non-descriptive content, it would be semantically inappropriate for me to use the informal pronoun when our relationship is more formal. But the impropriety is not misdescription.

¹⁹ Again, the distinction is not between kinds of content, but between kinds of language – language that describes and language that does not.

Consider also ‘is a Boche,’ a derogatory term for Germans.²⁰ As a predicate, its descriptive extension is the class of Germans, and it is used to describe people as being German. In addition to this descriptive aspect to meaning, it is inappropriate to use this term in any context unless one thinks ill of Germans. However, in calling someone a Boche one is not describing or representing that one has this negative attitude. So thinking ill of Germans, or something like it, is a non-descriptive use condition of the term.

I mention these other bits of language to warm us up to the idea that terms can have contents that are not described. The proposal is that normative terms are governed by attitude-involving w-rules that fix *non-descriptive* content. If we combine this with descriptive maximal vagueness, we get what I call *GAP Non-cognitivism*: i) thin normative terms are maximally vague with respect to descriptive semantics, as fixed by f-rules, and ii) they have non-descriptive use conditions fixed by w-rules that are sensitive to the non-cognitive attitudes of the speaker. GAP Non-cognitivism is gappy in at least two ways. F-rules leave a maximal *gap of semantic optionality* when it comes to use, and that gap is partly filled in by governing *attitude predication* – what happens when one follows the w-rules.

When a term is governed by two use rules, we must combine them to generate an overall use rule, and so an overall verdict on the semantic propriety of any use. Here is a proposal. Use of a term is overall semantically forbidden (inappropriate) iff any of the use rules forbid that use. Use of a term is overall semantically required iff any of the rules require that use.²¹ And use of a term is merely permissible iff none of the use rules require or prohibit its use. If the descriptive use rules neither require nor prohibit any use, as is the running hypothesis for thin normative terms, then the non-descriptive use rules will determine whether any given assertoric use is forbidden, required, or merely permissible.

²⁰ I follow Williamson (2009), who follows Dummett, in using this term. No offense intended.

²¹ If a use is forbidden by one rule and required by another, presumably we have an incoherent meaning.

GAP Non-cognitivism satisfies the two non-cognitivist hypotheses. It makes thin normative terms descriptively contentless and it explains the action-guiding qualities of the discourse just as speaker subjectivism does, the only difference being that GAP Non-cognitivism posits attitude-sensitive w-rules rather than attitude sensitive f-rules. It can also handle Frege-Geach problems. Let me explain.

4.2 Frege-Geach Payoff

Return to this argument one last time.

1. If tormenting the cat is wrong, then tormenting the dog is wrong.
2. Tormenting the cat is wrong.
- C. Tormenting the dog is wrong.

To explain the bundle of logical, semantic, and inferential properties associated with Frege-Geach problems, the classical statement of validity—if the premises are true the conclusion must be true—needs clarification to accommodate non-descriptive aspects to meaning. So let us say an argument is valid just in case, within any given context, if the use conditions of the premises are satisfied, then the use conditions of the conclusion must be satisfied. Basically, we look at all interpretations (assignments of use conditions for non-logical terms) where the use conditions are satisfied for the premises, and see if the use conditions of the conclusion are also satisfied. Satisfaction of use conditions here is akin to being true under the classical conceptions of validity, but it explicitly includes descriptive semantics and non-descriptive semantics. As with classical and vague-theoretic validity, where it does not matter what truth values the premises actually take, here it does not matter whether the use conditions of the premises are actually satisfied or not. When validity is understood as satisfied use-condition preservation, the argument is valid. There is no way of assigning use conditions to the non-logical vocabulary of (1) and (2) such that, within a single context of use, those use conditions can go satisfied without the use conditions of (C) going satisfied.

We can further say that one makes a logical mistake in accepting the premises and denying the conclusion in the same context, for the use conditions of (1), (2) and (\sim C) are not jointly satisfiable in the same context of use. What matters is whether, within a context, there is a way to satisfy the use conditions of the premises without satisfying the use conditions of the conclusion.²²

Concerning justification, if one is justified in accepting the premises and has no greater justification for rejecting the conclusion, one has justification for accepting the conclusion. We saw that maximal vagueness in descriptive semantics did not handle this property very well. It was unclear at best why many cases appear to be far from any borderline. And though it was permissible to predicate of these putatively borderline cases, it was hard to see why anyone could take themselves to have reason to do so.

GAP Non-cognitivism helps. When we aggregate satisfaction of descriptive and non-descriptive use conditions as suggested above, a sentence like “Tormenting the cat is wrong” comes out overall satisfied and semantically required if the speaker has the requisite negative attitude, despite the fact that ‘wrong’ is descriptively maximally vague. This is because the overall use conditions are satisfied iff no set of use conditions go unsatisfied and some set go satisfied. For ‘wrong’ the descriptive use conditions never go unsatisfied (for it has no anti-extension), and the non-descriptive use conditions sometimes go satisfied. Further, one justifiably accepts the sentence whenever one has good evidence that the use conditions are

²² Let me briefly take up the issue of penumbral connections mentioned in note 10. The non-descriptive use rule says not to predicate ‘is wrong’ of an action that has a property of which the speaker disapproves. Suppose Jack says that causing the dog n units of pain is wrong, but causing the dog $n+10$ units of pain is not wrong. This looks fairly similar to the standard penumbral connection case, where one predicates ‘bald’ of one person, and ‘not bald’ of a person with less hair. We can explain the similarity, for it is not easy to attribute to Jack a sensible psychology *and* competence with the use rule for ‘wrong’. If he follows the use rule in calling an action that causes n units of pain wrong, what could he disapprove of in action that would make it semantically appropriate for him to call causing $n+10$ amount of pain not wrong? It is not *impossible* to have such a psychology, but if Jack has a sensible psychology he violates penumbral connections with these predications.

satisfied. In the case of ‘wrong’, this would typically be evidence that the action has some feature one disapproves of in action. In this way maximal descriptive vagueness need not stand in the way of justifiably accepting a sentence. Evidence that the use conditions are in place is all one needs. One does not need descriptive content.

What of logical composition? As you might guess, we simply compose contents. This amounts to composing use conditions, which include descriptive conditions and non-descriptive conditions. So, “Either normative terms are maximally vague or tormenting the cat is wrong” is an output of a function on the use conditions (contents) of “Normative terms are maximally vague” and “Tormenting the cat is wrong.” The use conditions of the disjunction are satisfied iff the use conditions of first disjunct are satisfied (that is, normative terms are maximally vague) or the use conditions of the second disjunct are satisfied (that is, tormenting the cat has some property you all-in disapprove of). So, clearly, one can assert this disjunct in semantically appropriate fashion even if one lacks any negative attitudes toward Tormenting the cat, so long as the first disjunct’s use conditions are satisfied.

On this view normative sentences compose exactly as they would under (sophisticated) speaker subjectivism. For the difference between GAP Non-cognitivism and speaker subjectivism is not a difference in contents assigned, but rather a difference in what the language is doing with those contents. Under one hypothesis contents are described; under the other they are not. In the contexts we are considering, logical operators only care about contents, not whether the language is describing them or not.

Now, on this view “Tormenting the cat is wrong” is not used to describe its content, but “Tormenting the cat does not cause pain” is used to describe its content. What about “Either tormenting the cat is wrong, or it does not cause pain”? Is its overall content described or not? The best thing to say is that its overall content is a mixture of descriptive content and non-descriptive content. The complex sentences is correctly assertable iff one of

its disjuncts is correctly assertable, where the assertion of one disjunct would describe, and the assertion of another would not. We know what conditions must be in place for each disjunct to be correctly assertable. And we know the rules one grasps when one grasps the meaning of the complex sentence. If someone asserts the disjunction you can bet that they believe that tormenting the cat does not cause pain or tormenting the cat has some property they disapprove of in action.²³

4.3 To Kick Away the Descriptive Side?

We are at a crossroads. We might not need a descriptive semantics of maximal vagueness to solve Frege-Geach problems. For the explanations of Frege-Geach properties given above apply just as well to what we can call WUR Non-cognitivism, or *when use rule non-cognitivism*. WUR Non-cognitivism differs from GAP Non-cognitivism in dropping the descriptive dimension to meaning for thin normative terms. That is, it does not posit that thin normative terms are governed by f-rules at all, not even ones that deliver maximal vagueness. It only posits w-rules and associated non-descriptive content. Having content suffices to explain logical, semantic, and inferential roles, after all, so it is not clear that we need an extra layer of f-rules and maximal vagueness.

Whether a non-cognitivist *needs* a descriptive semantics of maximal vagueness is one thing. Whether it is *desirable* is quite another. So let me mention some reasons one might want to keep maximal descriptive vagueness in the picture.

First, some pure non-cognitivists prefer a pragmatic theory of attitude expression. Such non-cognitivists would need something like maximal vagueness on the descriptive side to get the semantic, logical and inferential roles we need.

Second, you might be independently convinced that there is a descriptive semantics for thin normative terms. Perhaps you like causal regulation as a meta-semantic theory for

²³ Future work would have to address normative thoughts more directly. Here, the focus is on language.

how descriptive content gets fixed. A pure non-cognitivist can then keep descriptive-content-fixing rules as devices for modeling the descriptive-content-fixing facts (here, causal facts), while maintaining that they are woolly enough to fix no content.

Third, maximally vague descriptive semantics still has a role to play if having a descriptive dimension to meaning explains why a term is capable of taking the predicate place. Recall that many of our terms that ostensibly have only non-descriptive meanings, like ‘ouch’, do not take the predicate place. That might be because they have no descriptive semantics – not even the façade of maximal vagueness. If so, pure non-cognitivists will want to retain maximal descriptive vagueness to explain why their descriptively contentless terms are predicates.

We now have several interesting options for non-cognitivists to consider: 1) maximal descriptive vagueness and pragmatic action-guidance (VAP Non-cognitivism), 2) maximal descriptive vagueness plus semantic action-guidance via non-descriptive w-rules (GAP Non-cognitivism), 3) no descriptive semantics, only semantic action-guidance via non-descriptive w-rules (WUR Non-cognitivism), and the non-descriptive content they help to fix. I prefer (2) and (3) over (1), but I will not argue their merits any farther.

5 Future Directions: Truth vs Semantic Propriety; Disagreement; Motivation; Why Language that Works Like This?

Of course, a complete metaethic must explain many other things. My main aim here is to suggest alternative locations for waging the battles. But I want to close the discussion with some suggestive remarks about what these versions of non-cognitivism have to say about other bugbears.

First, truth. There are two things I want to think more about. For one, is the concept of truth really the concept we want to employ in theoretical semantics? My current thinking is that it is a mistake to employ talk of truth in semantic theory to begin with. We should be

employing a theoretical term—something like *semantic correctness* relative to a context and circumstance of evaluation—so it is not confused with an ordinary language term with an arguably minimalist meaning given by instances of T-schema.²⁴ We might be able to use the semantic theory and its posits to explain the meaning of ‘true’, to be sure, but we do better to employ ‘semantic correctness’ in the theory itself.

If we do this we are then able to make some important distinctions. For one, we can acknowledge when someone uses language in a semantically appropriate fashion, but otherwise take issue with what they say or think by saying things like “that’s not true”. Semantic propriety is one thing, truth another. To illustrate, suppose that tormenting the cat exhibits some property Able disapproves of in action, but fails to exhibit some property that Bea disapproves of in action. Able will want to distinguish the *semantic propriety* of Bea’s assertion “Tormenting the cat is not wrong” from the *truth* of “Tormenting the cat is not wrong”. If we speak of semantic correctness and cognates in semantic theory we are free to make this desirable distinction.

The second thing I want to thin more about is the relationship between truth and accurate description. It can be easy to conflate them, but I think we should resist doing so. One attractive option is *minimalism* about truth and falsity, but non-minimalism about representational accuracy. Perhaps all we can say about truth is exhausted by the T-schema and its instances. S is true iff S. “Tormenting the cat is wrong” is true if and only if tormenting the cat is wrong. When it comes to representational accuracy, non-cognitivism will say that my acceptance of “Tormenting the cat is wrong” does not have representational content. In terms of the resources employed here, we can cash this out in terms of *when* use rules vs. *for* use rules, non-descriptive content vs. descriptive content, etc. We will not be minimalist about some of these theoretical notions. Thanks to minimalism about truth, my

²⁴ See Field (1994, p. 409-22) for a discussion of disassociating truth with the other semantic notions we used to theorize about vagueness.

acceptance of the truth of that sentence is both a free move and it commits me to no more representational purport than the original acceptance. Alas, further consideration is for another time.²⁵

Here is an obviously related matter: disagreement. Suppose Able says “Tormenting the cat is wrong”, and Bea replies “Tormenting the cat is not wrong”. It is a datum that they disagree. On GAP Non-cognitivism and WUR Non-cognitivism, Able’s and Bea’s assertions could come about in the following way. Able says what she says because she all-in disapproves of properties F1, F2 and F3 in action, and she thinks tormenting the cat exhibits one of those properties. Bea says what she says because she all-in disapproves of properties G1, G2 and G3 in action, and she thinks tormenting the cat does not exhibit any of those properties.

Though the content of what they say is not inconsistent, I want to mention two possibilities for respecting their disagreement. One option is to hold out for disagreement in attitude. W-rules help with this. They help to assign content *without* putting the presence of our attitudes, or the properties that are their objects, *at issue* in conversation as accurate representations of how things are. Though our approach assigns contents in a speaker-variant way, at the same time it does not imbue those things with representational purport, and so denials, accusations of falsity and the like are not going to contest *a way things have been described*. Indeed, it is fair to say that one of the main points of having a w-rule, rather than a f-rule, is to avoid putting content at issue as an accurate representation of how things are. This opens the door for handling any normative disagreement as disagreement in attitude.

A second option is to go minimalist about disagreement. Given the tight connection between disagreement and ascriptions of falsity, and agreement and ascriptions of truth, and given the attractions of minimalism about truth, minimalism about disagreement should be

²⁵ Another thing to address is how contents are fixed in intensional contexts.

appealing. There, the thought is that there is nothing more to disagreement about S other than one thinking S is true and the other that it is false. Just as we might jettison a correspondence theory of truth, we might jettison a theory of disagreement as incompatible representings. And we might say there is nothing more for a theory of disagreement to explain that requires more than the minimalist theory. Again, additional discussion is for another day.

Here is one last comment about GAP Non-cognitivism and WUR Non-cognitivism, our two theories that make the action-guiding element semantic. It is an asset of our semantic theory of action-guidance that it delivers a weak form of motivational internalism. For us, accepting the sentence “It would be wrong for me to step on gouty toes” does not constitutively involve non-cognitive attitude. Rather, having such an attitude is part of the semantic correctness conditions for accepting the sentence. For one to correctly accept it, tormenting the cat must have some property one all-in disapproves of in action. As with any semantic correctness conditions, one can accept the sentence while unwittingly violating its correctness conditions. That is, we can think an action is wrong without knowing what features it has that we find objectionable. In that case, we would not expect a very strong motivational profile, but a profile more akin to beliefs that there is *something* about the action that we would find objectionable. Perhaps typically we are aware of the (believed) features that we find objectionable and so typically when one accepts a first-person normative sentence one will be motivated accordingly. But on the present view this strong motivation is not *necessarily* for all normative judgments. This gets us a special connection to motivation that standard non-normative thoughts lack, but not any stronger, highly controversial, motivational internalist theses. In particular, we do not have the view that thinking x is wrong is partly constituted by a conative attitude toward x or some of its (believed) properties. These are provisional remarks, of course, and further consideration of motivational internalism is for another time.

Even if we have seen some non-cognitivist options with novel explanatory resources, you might wonder about the ground on which they are built. *Why in the world* would we have maximally vague terms? Why would we have *when* use rules, as opposed to *for* use rules, for any term?

This requires a different sort of explanation of the language, not how it works but why we have it. Other non-cognitivists have already provided part of the answer. Social beings with a plurality of differing and conflicting non-cognitive attitudes need a way of communicating their non-cognitive attitudes to others. They need to indicate what they are for and against, and to discern what others are for and against. Further, if coordination and cooperation are to get off the ground, such beings need a way to reason through what to be for and against, both individually and together (c.f. Gibbard 2003, p. 13).

Just *how* we do this—how we can think and talk about what attitudes to have and what to do without putting at issue the nature of our attitudes, or their contents—has proven elusive. The resources here help fill the bill. Granted, it would be nice were there some non-normative fragment of language that clearly works just as I have hypothesized thin normative language to work. I could then just draw analogies, and we would already be sympathetic to the ingredients of the theory. Unfortunately, I know of no clear examples that can aid digestion of our three hypotheses in this way. But one of the dominant themes in the history of metaethics is that there seems to be something *sui generis* about normativity. So perhaps the absence of a meaning homologue is not so bad. Suffice it to say that we have started to take non-cognitivism down new roads, thanks to vagueness and use-conditional meaning.*

* [Acknowledgments].

Bibliography

- Åqvist, L. (1964). Vagueness and Value. *Ratio*, 6(2), 121–127.
- Baker, D., & Woods, J. (2015). How Expressivists Can and Should Explain Inconsistency. *Ethics*, 125(2), 391-424.
- Blackburn, S. (1984). *Spreading the Word: Groundings in the Philosophy of Language*. Oxford Oxfordshire; New York: Clarendon Press; Oxford University Press.
- Brandt, R. B. (1954). *Hopi Ethics: a Theoretical Analysis*. Chicago: Chicago University Press.
- Charlow, N. (2013). The problem with the Frege–Geach problem. *Philosophical Studies*, 167(3), 635–665.
- Chrisman, M. (2012). On the Meaning of “Ought.” In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Volume 7* (pp. 304-32). Oxford: Oxford University Press.
- Cobreros, P. (2011). Supervaluationism and Classical Logic. In R. Nouwen, R. van Rooij, U. Sauerland, and H.C. Schmitz (Eds.), *Vagueness in Communication* (pp. 51–63). Berlin, Heidelberg: Springer-Verlag.
- Dreier, J. (1996). Expressivist Embeddings and Minimalist Truth. *Philosophical Studies*, 83(1), 29–51.
- Dreier, J. (1999). Transforming Expressivism. *Nous*, 33(4), 558–572.
- Dreier, J. (1990). Internalism and Speaker Relativism. *Ethics*, 101(1), 6–26.
- Field, H. (1994). Disquotational Truth and Factually Defective Discourse. *The Philosophical Review*, 103(3), 405-52.
- Fine, K. (1975). Vagueness, Truth and Logic. *Synthese: An International Journal for Epistemology, Methodology and Philosophy of Science* 30: 265-300.
- Geach, P. (1965). Assertion. *The Philosophical Review*, 74(4) 449-465.

- Geach, P. (1960). Ascriptivism. *The Philosophical Review*, 69(2), 221–225.
- Gert, J. (2012). Basic Normative Terms. In *Normative Bedrock* (pp. 34–70). Oxford: Oxford University Press.
- Gert, J. (2007). Cognitivism, Expressivism, and Agreement in Response. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Volume 2* (pp. 77-110). Oxford: Oxford University Press.
- Gibbard, A. (2003). *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.
- Gutzmann, D., & Gärtner, H.-M. (2013). *Beyond Expressives: Explorations in Use-Conditional Meaning, Current Research in the Semantics/Pragmatics Interface (CRiSPI) 28*. Leiden: Brill.
- Horgan, T. & Timmons, M. (2009). Expressivism and Contrary-Forming Negation. *Noûs-Supplement: Philosophical Issues*, 19, 92-112.
- Horgan, T. & Timmons, M. (2006). Cognitivist Expressivism. In T. Horgan and M. Timmons, (Eds.), *Metaethics after Moore* (pp. 255–298). Oxford: Oxford University Press.
- Horgan, T., & Timmons, M. (1992). Troubles on Moral Twin Earth: Moral Queerness Revived. *Synthese*, 92(2), 221–260.
- Jackson, F., & Pettit, P. (1998). A Problem for Expressivism. *Analysis*, 58(4), 239–251.
- Price, H., Blackburn, S., Brandom, R., Horwich, P., & Williams, M. (2013). *Expressivism, Pragmatism and Representationalism*. Cambridge University Press.
- Ridge, M. (2014). *Impassioned Belief*. Oxford University Press.
- Rosen, G. (1998). *Blackburn's Essays in Quasi-Realism* (New York: Oxford University Press).

- Sainsbury, M. (1990). Concepts Without Boundaries. Inaugural Lecture, Department of Philosophy, King's College, London; reprinted in R. Keefe and P. Smith (Eds.), *Vagueness. A Reader* (pp. 251-264), Cambridge (MA): MIT Press, 1996.
- Schiffer, S. (2002). Moral Realism and Indeterminacy. *Philosophical Issues*, 12(1), 286–304.
- Schroeder, Mark. (2008a). *Being for: Evaluating the Semantic Program of Expressivism*. Oxford, UK; New York: Clarendon Press; Oxford University Press.
- Schroeder, M. (2008b). What is the Frege-Geach Problem? *Philosophy Compass*, 3(4), 703–720.
- Schwartz, J., & Hom, C. (forthcoming). Why the Negation Problem Is Not a Problem for Expressivism. *Noûs*.
- Searle, J. R. (1969). *Speech Acts*. Cambridge University Press.
- Shapiro, S. (2006). *Vagueness in Context*. Oxford England; New York: Clarendon Press; Oxford University Press.
- Smith, M., & Stoljar, D. (2003). Is There a Lockean Argument against Expressivism? *Analysis*, 63(1), 76–86.
- Sorensen, R. A. (1990). Vagueness Implies Cognitivism. *American Philosophical Quarterly*, 27(1), 1–14.
- van Roojen, M. (1996). Expressivism and Irrationality. *Philosophical Review*, 105(3), 311-335.
- Williamson, T. (2009). Reference, Inference, and the Semantics of Pejoratives. In J. Almog and Po Leonardi (Eds.) *The Philosophy of David Kaplan* (pp. 137-159). Oxford University Press.