

Copyright

by

Taylor-Grey E. Miller

2018

**The Report Committee for Taylor-Grey E. Miller certifies that this is the approved
version of the following report:**

In What Sense (If Any) Is Meaning Normative?

**APPROVED BY
SUPERVISING COMMITTEE:**

Lawrence Buchanan, Supervisor

Cory Juhl

In What Sense (If Any) Is Meaning Normative?

by

Taylor-Grey Miller, B.A.

Report

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

MASTER OF ARTS

The University of Texas at Austin

August 2018

Abstract

In What Sense (If Any) Is Meaning Normative?

Taylor-Grey Miller, M.A.

The University of Texas at Austin, 2018

Supervisor: Lawrence Buchanan

“Meaning is normative” has become a popular dictum in the philosophy of language. Moreover, some proponents of the dictum have appealed to the normativity of meaning in order to demonstrate the hopelessness of a reduction of semantic to non-semantic vocabulary. The upshot, some maintain is that this stands in the way of a naturalistically respectable theory of meaning. The aim of this paper is to get clear on what exactly the dictum should amount to in order to pose a problem for a naturalistic theory of meaning and then see whether the dictum is true. I will argue that the most plausible version of the thesis that meaning is normative poses no real threat to the possibility of naturalism about meaning.

Table of Contents

List of Tables	vi
Chapter 1: In What Sense (If any) Is Meaning Normative?	1
1.1 Introduction	1
1.2 Preliminaries.....	2
1.2.1 What a Theory of Meaning is a Theory of.....	2
1.2.2 First Pass as the Normativist Thesis.....	3
1.2.3 A Change of Tact.....	12
1.3 The Normative Status of Semantic Rules.....	16
1.3.1 Semantic Rules as Constitutive Rules	16
1.3.2 Norms of Execution are Instrumental Norms.....	20
1.3.3 Defending P1	24
1.3.4 Defending P2: Semantic Rules Do Not Specify Victory States.....	29
1.4 A Worry	34
1.5 Preliminary Conclusion	36
1.6 Articulating the Meaning Norms	36
1.6.1 Logical Space.....	37
1.6.2 Harman's Worries and Their Linguistic Counterparts	43
1.7 Conclusion.....	52
Works Cited.....	53

List of Tables

Table 1:	Bridge Principles:	40
----------	--------------------------	----

Chapter 1

In What Sense (If Any) Is Meaning Normative?

1.1 Introduction

“Meaning is normative” has become a popular dictum in the philosophy of language. Moreover, some proponents of the dictum have appealed to the normativity of meaning in order to demonstrate the hopelessness of a reduction of semantic to non-semantic vocabulary. The upshot, some maintain is that this stands in the way of a naturalistically respectable theory of meaning. The aim of this paper is to get clear on what exactly the dictum should amount to in order to pose a problem for a naturalistic theory of meaning and then see whether the dictum is true. I will argue that the most plausible version of the thesis that meaning is normative poses no real threat to the possibility of naturalism about meaning.

In section 2, I present a standard formulation of the normativist thesis and present some standard as well as novel arguments why it fails. I further argue that the thesis can be rescued by understanding not as a semantic but rather a meta-semantic thesis; that the meaning of expressions are grounded in semantic rules governing their use. In section 3, I present the commitments of the meta-semantic version of normativism and argue that semantic rules are not normative in the way needed to impose a substantive constraint on an acceptable theory of meaning. In section 4, I consider and respond to an objection to my argument in section 3. In section 5, I argue that these results generalize in important ways. Section 6, presents a systematic discussion of the norms of language use many have taken to follow from an expression’s

meaning. I give novel considerations in favor of skepticism about any special normative relation between an expression's meaning and how speakers ought to use it. This underscores the results of the first part of the paper. Given that there is no plausible meaning norm, it follows that the normativity of meaning poses no real threat to the possibility of naturalism about meaning.

1.2 Preliminaries

1.2.1 What a Theory of Meaning is a Theory of

The phrase “theory of meaning” is ambiguous. We need to distinguish between:

- (1) A theory that, for some language (say Chinese), produces a specification of meaning for each well-formed sentence in that language.
- (2) A theory which gives an analysis of meaning in terms of other concepts, or at least situates the concept of meaning with respect to other concepts.

Theories of the first kind aim to produce theorems of the form *Sentence S in Language L has meaning P*. Usually as a simplifying assumption, theorists aiming to give a meaning theory of this type regard languages as static unchangeable things.¹ The upshot of this simplifying assumption is that for example, the language which is in fact currently spoken in mainland China has the non-contingent property that in it “下雨” means that it is raining. A language in which those expressions did not mean what they do would be a different language. What is taken to be contingent is the fact that a certain population

¹ See for example, Davis (1981), Gamut (1991), and Kamp (2008). Or any introductory formal semantics textbook for that matter.

(e.g., Mainland Chinese) speak the language they do (e.g., Mandarin). If things had gone differently (say) if the Huns hadn't successfully invaded when they did, a very different language could have been spoken in China. Ignoring for the sake of simplicity, contextual sensitivity and ambiguity, we could, following Lewis, consider a language L an ordered pair $\langle S, M \rangle$ where S is a set of sentences and M is a set of meaning specifications, one for each sentence in the set S.

If $\langle S, M \rangle = L$ and $s \in S$ for which M specifies:

s means that P

then s is a sentence of L and s means P in L.²

Thus, theories₁ of meaning are semantic theories. What becomes of immediate philosophical interest is giving an account of the conditions under which a semantic theory is the correct semantic theory for speakers of a particular language. In other words, giving an account of the facts in virtue of which a language L is the *actual* language of a certain population. To state these conditions would be to give a theory₂ of meaning, that is a philosophical elucidation of meaning. Thus, to give a theory₂ of meaning is to give a meta-semantic theory.

1.2.2 First Pass at the Normativist Thesis

Both proponents and opponents of the normativity of meaning agree on the following: the right meaning theory for a language specifies the correctness conditions for expressions in that language. This is supposed to demarcate meaningful language use from mere noise-making.³ For example, if we take

² Lewis (1975, 3)

³ Hattiangadi (2006), (Boghossian (1989), and Gluer (1999)

our language to be English and the relevant expression to be “chair” then the right meaning theory of English will specify the set of conditions that are necessary and sufficient for the expression “chair” to be true of the members of its extension. We can represent this disquotationally, as in:

For any object x , “chair” correctly applies to x iff x is a chair

We can further generalize and represent the correctness conditions of any expression by the following formula. Where e is an expression, F gives its meaning, and f is the feature or set of features in virtue of which F applies, we get:

Correctness Conditions (CC): e means $F \rightarrow \forall x (e \text{ correctly applies to } x \text{ iff } x \text{ is } f)$ ⁴

The normativist and anti-normativist then can agree on all the relevant semantic facts. They can share the same theory₁ of meaning. The standard point of departure for the normativists is in what they take to follow from the relevant semantic facts. They hold that the semantic facts and therefore the correctness conditions of expressions in the language impinge on speakers in normatively significant ways. Facts about the correctness conditions for the expression “chair” give rise to norms regarding how the expression “chair” ought to be used. This has quite a bit of intuitive appeal. Often, we cite semantic facts as explanations for why speakers ought to use expressions in certain ways. In fact, I encountered such a case recently.

⁴ Some might object that the simplifying assumption in the previous section is not harmless. That may be true, but it’s not dialectically harmful. Proponents of the normativity thesis go in for this simplifying assumption when they state what ought to be considered dialectically uncontroversial e.g., the correctness conditions.

At a family gathering, one of the older children, call him John, read a book to the younger children. When John got to a part of the story that featured zebras, one of the younger children, call him Ben, asked him, “what are those things supposed to be called?” John responded “zebras!” while pointing to the animals. Puzzled, Ben asked why, to which John responded “Because that’s what ‘zebra’ means.” Ben took that as a satisfying explanation and they moved on with the story. Examples like this can be multiplied and they all seem to support the normativist’s claim. Additionally, it seems plausible that when we hear novel expressions, we tacitly take the speaker to utter that expression *because* it has the meaning that it does, and we often exploit this to extrapolate the meaning of the novel expression from the circumstances of its usage. In light of these considerations some normativists have given different versions of the following argument. Call it the Face-value Argument:

1. e means $F \rightarrow \forall x (e \text{ correctly applies to } x \leftrightarrow x \text{ is } f)$
2. $\forall x (e \text{ correctly applies to } x \leftrightarrow x \text{ is } f) \rightarrow \text{speaker S ought to (apply } e \text{ to } x \leftrightarrow x \text{ is } f)$
- \therefore 3. e means $F \rightarrow \text{speaker S ought to (apply } e \text{ to } x \leftrightarrow x \text{ is } f)$

An expression’s meaning, they argue, straightforwardly entails norms regarding its use. This is supposed to be a problem for naturalistic accounts of meaning because no such reductive account could do justice to this data. To sharpen the point, take for example, Kripke’s famous case of “+”.⁵ By substituting in appropriately to the above argument schema we can derive:

⁵ Kripke (1982)

“+” means plus \rightarrow one *ought* to apply “+” to the triple $\langle 57, 68, 125 \rangle$

Any adequate analysis of the antecedent meaning fact would have to do justice to this normative entailment. And as the story goes, there isn't much room for the naturalist to maneuver. No behavioristic or mentalistic account can meet this constraint, so no such account can be correct. For example, one might be tempted think that when a person means PLUS by “+”, he does so in virtue of being disposed to assent to certain sums, e.g. “ $1 + 2 = 3$ ” and not others, e.g. “ $1 + 2 = 4$ ”, but surely no such facts about dispositions could possibly explain what the person *ought* to accept! This would be tantamount to deriving an ought from an is. How could it provide the requisite justification? How could it guide me to my answer? How could it be that I answer in light of what “+” means? Call this position face-value meaning normativism.

Face-value Meaning Normativism: Naturalistic accounts of meaning are ruled out by the fact that what an expression means immediately implies a statement about what speakers ought to do with that expression.

This position has been advocated by a number of philosophers.⁶ However, despite the many notable adherents, it false.

The reason why the face-value argument fails was brought out most forcefully by Paul Horwich. It is not enough to motivate a non-naturalism to merely point out that what we are trying to give a theory of has normative import. He puts it in the following way:

⁶ See e.g., Blackburn (1984), Bloor (1997), Brandom (1994), Gibbard (1994), Miller (1998), Glock (1996), McDowell (1994), and most recently by Daniel Whiting (2009).

It is fairly clear that something can perfectly well have normative import without being constitutionally normative. Surely, infection with smallpox is a bad thing—something one ought to try to prevent—even if (as is plausible) our concept of it is characterizable in purely biological terms. Similarly, there would seem to be no incoherence in fundamentally valuing the preservation of giant pandas, even though the existence of these animals is in itself a non-evaluative state of affairs. These examples illustrate what is surely a common propositional structure: the subject identifies something in ‘descriptive’ (entirely non-evaluative) terms; and the predicate proceeds to appraise it, to assess it, and to evaluate the identified phenomenon. Thus, it is evident that the issue of whether [meaning is] intrinsically normative is not settled simply by calling attention to [meaning’s] evaluative implications.⁷

The moral is that something’s having normative implications provides no reason at all to suspect that it is intrinsically normative. Killing is *prima facie* wrong, however, it is totally uncontroversial that we can characterize the act of killing in purely descriptive terms. Moreover, it straightforwardly follows from “x kills y” that x did something he ought not to have done. So, if the normative import of the property “is such that he killed y” does not determine whether the property is intrinsically normative, it follows that the normative import of a meaning-property isn’t enough to guarantee that that it is intrinsically normative. And since we can give an account of killing in wholly descriptive terms, why not think by the argument’s lights, the same could be

⁷ Horwich (2005, 109)

said for meaning?⁸ Normativists then are not vindicated merely by the fact that meaning has certain normative implications.

This brings out a deep tension for the face-value normativist. The view seems well motivated, but the face-value argument struggles to explain why the intuitive motivations should introduce real worries for naturalism in a non-question begging way. For any way of running the face-value argument, the naturalist will always have a way of resisting the conclusion. To see why this is so in greater detail, it will be useful to look at some contemporary face-value responses to the line of defense Horwich gives. We will see that they either beg the question against the naturalist, or else run afoul of equivocating on some of the central notions that feature in the argument.

The most recent and most comprehensive defense of face-value normativism has been offered by Daniel Whiting.⁹ He argues that he can give principled reasons for why meaning is intrinsically normative and it has to do with an important difference between the kinds of normative implications meaning has. Whiting claims that there is a dis-analogy between paradigmatic instrumental or means-end obligations and semantic obligations. In the case of paradigmatic instrumental obligations, the “ought”s are contingent on the desires: if your desires change, so does what you ought to do. For example, suppose that you want to go for a walk in the snow, but you don’t want your ears to freeze. Given these desires, you ought to (say) wear a beanie. However, if you change your mind, and decide that you don’t mind having cold ears, it is no longer true that you ought to wear a beanie. If you do wear a beanie, you

⁸ Horwich (2005, 110)

⁹ Whiting (2007)

will not have done something you ought not to have done. Whiting argues that in the semantic case, the “ought”s are not contingent on desires. He says:

In contrast, given what “rich” means, that I ought to apply the term to a person only if she is rich does not seem contingent upon (say) my desire to speak truthfully. If that desire changes, and I apply the term to a poor person, it remains the case that I am not applying it as it should be applied, but rather incorrectly. Here, it seems one is properly entitled and it makes full sense to judge that, desire notwithstanding, I am using the expression wrongly.¹⁰

Whiting claims that whether I ought to apply a term is not contingent on my desire to speak truthfully because, desire notwithstanding, it remains the case that:¹¹

- (i) an application of “rich” to a poor person is *incorrect*,
- (ii) an application of “rich” to a poor person is *wrong*, and
- (iii) “rich” *should not* be applied to a poor person.

Whiting suggests that (i), (ii) and (iii) have implications for whether I ought to apply the term “rich” in any given situation. He claims that (i), (ii) and (iii) are not contingent on the desire to speak truthfully and concludes that how I ought to use a term is not contingent on my desires (or other, non-semantic considerations). However, as Anandi Hattiangadi has pointed out, this argument either equivocates, or begs the question.¹² All of the italicized terms

¹⁰ Whiting (2007, 139)

¹¹ Whiting (2007, 139)

¹² Hattiangadi (2008)

in (i), (ii) and (iii) have both normative and non-normative interpretations. If we assume a normative interpretation of these key terms, the argument becomes question-begging. If we assume a non-normative interpretation of the key terms, then (i), (ii) and (iii) do not have normative implications of the sort Whiting claims and the argument is invalid.

First, consider (i). Recall that for both normativists and naturalists all instances of CC are platitudinous. From this it follows that:

“rich” applies correctly to x if and only if x is rich.

From this it follows that if you apply “rich” to someone who is poor, then you have applied “rich” incorrectly. However, contexts like these it is quite unclear whether, ‘correct’ and ‘incorrect’ function as normative terms. To say that you have applied “rich” incorrectly is very plausibly just to say that you have applied “rich” to someone of whom “rich” is not true.¹³ Since that does not imply anything about what you ought to say, the fact that (i) is not contingent on your desires does not show that you have a noncontingent obligation not to apply “rich” to a poor person.

The same goes for (ii). Both “right” and “wrong” have non-normative uses. For example, to give the right answer to a question, or in an exam, is to give the answer that satisfies the expectations of the questioner or examiner. To say that an answer is right in this sense is not to say that it is the answer that you ought to give. If “right answer” just meant “answer that you ought to give,” it would sound odd to say “you should not give the right answer,” or “you should give the wrong answer.”¹⁴ However, these are perfectly

¹³ Hattiangadi (2008, 60–62)

¹⁴ Hattiangadi (2008, 62–63)

reasonable things to say in some situations; for example, if by answering to satisfy the examiner's expectations you will score too many points and throw off the class curve. Similarly, to say that your application of "rich" was wrong can be just to say that you have applied "rich" to someone of whom "rich" is not true. And the fact that you have applied "rich" to someone of whom "rich" is not true does not imply that you have done something you ought not to have done. Thus, from the fact that (ii) is not contingent on your desire to express a truth, it does not follow that you have a non-contingent obligation not to apply "rich" to a poor person

Similar remarks go for (iii). Beanies ought to be used to stay warm in the snow. Suppose that I do not mind getting cold and go out in the rain without a beanie. This does not alter the fact that beanies ought to be used to stay warm in the snow. Even if I wish to use a beanie to scoop up snow to make a snow man, it remains a beanie—something which ought to be used to stay warm in the snow. By using a beanie to scoop up snow to make a snow man, I will not be using it in the way that it ought to be used; I will be using it *incorrectly*, even *wrongly*. Yet, we certainly wouldn't want to say that a beanie's function is intrinsically normative. By the same token, if I apply "rich" to a poor person, I will not be using it in the way that it ought to be used; I will be using it *incorrectly*, even *wrongly*. Why think that by this reasoning we should think that meaning is intrinsically normative?

Again, we see that the naturalist has a way out. The face value argument simply lacks the resources to compel the normativist conclusion. Thus, it fails. In light of the failure of the face-value argument, many have rejected the normativity thesis. However, I think there is something of interest that the normativist has to say, it just isn't something that is best captured by the face-

value approach. While it seems quite natural to present the view this way, normativists don't do themselves any favors arguing along these lines. So, the normativist needs to change tact.

1.2.3 A Change of Tact

While it seems plausible that facts about how expressions are to be used in some sense follow from semantic facts like CC, one question that naturally arises is "In virtue of what is it the case that facts about how I ought to speak follow from the relevant semantic facts?" A promising way of repairing the normativist thesis emerges when we formulate the view in terms of an answer to this question. The view should then be understood as claiming that the facts which ground the semantic facts are essentially normative facts. When giving an account of that in virtue of which (say) English is the *actual* language of a certain population, one must make appeal to essentially normative facts. The normativity then is supposed to emerge at the meta-semantic level. Call this meta-semantic normativism.

Meta-Semantic Normativism: the facts in virtue of which a semantic theory is the right semantic theory for a given population are essentially normative and thus rule out naturalistic theories of meaning.

This opens up a tempting place for the naturalist to dig in her heels. She may try to block the normativity from creeping in at the outset by rejecting the meta-semantic story the normativist offers. Many naturalists have pursued this line of argumentation.¹⁵ However, I don't think that the naturalist *has* to dig in her heels here. What I hope to show is that the naturalist can grant the

¹⁵ Boghossian, (1989), Hattiangadi (2007)

normativist her preferred meta-semantic account but deny that such an account really involves essentially normative facts. In other words, both parties can agree on what a theory₂ of meaning consists in but disagree on the nature of the relevant facts. This is a happy dialectical result since it allows the naturalist to engage with a maximally plausible and concessive version of the normativity thesis, and if the naturalist finds a way to resist the normativist's conclusion, the normativist will be quite hard pressed to find any way of further repairing the view.

What then is the normativist's preferred meta-semantic account? The single most widely held position is that the meaning of linguistic expressions is determined by or grounded in *rules for their use*.¹⁶ This position is most famously associated with Wittgenstein. He wrote: "Without these rules the word has as yet no meaning; and if we change the rules, it now has another meaning (or none), and in that case we may just as well change the word too"¹⁷ So, we can sharpen the normativist's meta-semantic thesis in the following way:

Semantic Rule-Following (SR): *e* means F just in case there is a semantic rule R governing the use of *e* in force for a given population.¹⁸

If we take on board the claim that meaning is grounded in semantic rules, a number of questions arise regarding the nature of the semantic rules. For

¹⁶ See e.g., Gluer (2009), Brandom (2000), Wilkforss (2001), and Kisselbach (2014).

¹⁷ Wittgenstein (1953, 133). This should come as no surprise given the centrality that rule-following has played in shaping the debate, no doubt in large part due to the role rule-following plays in Kripke (1982).

¹⁸ Gluer and Wilkforss (2018)

example, what kinds of rules must semantic rules be to play the role attributes to them? We will address this and related questions in the next section.

For now, we will focus on what it could mean to say that the semantic rules determining meaning are essentially normative. To get clear on what this could amount to, we will distinguish between two grades of normative involvement. Let a theory have grade 1 normative involvement just in case the theory *simpliciter* entails normative conclusions. This is the minimum degree of normative involvement that a theory₂ of meaning would need in order to threaten naturalism about meaning, since it is only at this level that a theory can give rise to non-instrumental norms. Let a theory have grade 2 normative involvement just in case the theory has normative consequences but only alongside other normative assumptions.

To borrow an example from Gillian Russell, in arithmetic, the fact that $67 + 58 = 125$ could entail that you shouldn't believe that you have 67 cents in one pocket, 58 cents in the other, and also believe that you have less than \$1.25 in your pockets. However, this isn't because $67 + 58 = 125$ is a normative fact, nor even because arithmetical facts have normative consequences; it's because if you have 67 cents in one pocket and 58 cents in the other it is false that you have less than \$1.25 in your pockets. In tandem with the non-arithmetical but normative fact that one shouldn't have false beliefs, this entails that you shouldn't believe you have less than \$1.25 in your pocket. So, the theory of arithmetic has normative consequences, but not on its own; only when taken alongside other albeit widely accepted normative assumptions.¹⁹

¹⁹ Russell (2017, 10)

If the semantic rules determining meaning exhibit only grade 2 normative involvement, then they are normatively inert. They have no normative consequences taken in isolation and give us only descriptive information regarding the relation between semantic and meta-semantic facts. As such it would be compatible with naturalism about meaning and could be normative in only the same sense as the arithmetic facts—merely instrumentally normative. Advocates of the normativity of meaning argue that semantic rules exhibit grade 1 normative involvement, and this is what is meant by the claim that meaning is essentially or intrinsically normative.

It is my contention that semantic rules determining meaning exhibit only grade 2 normative involvement. In conjunction with distinct but common normative commitments such as our communicative needs and interests, semantic rules have normative consequences. The upshot is, if this is the extent to which meaning is normative, then it shares this status with paradigmatically descriptive scientific theories, including those of physics and mathematics. Moreover, when these norms are widespread (as surely some kind of norms regarding communication are) the utterance of even a descriptive sentence will usually have normative implicatures. In many contexts, if you apply the expression “zebra” to something and I say, “But it doesn’t have stripes, and if it’s a zebra then it has stripes” you will rightly take me to be telling you more than can be extracted from the semantic content of the sentence alone. Often, an obvious consequence of what I said in the context of standard conversational norms is that the expression “zebra” was incorrectly applied, and so you ought not to apply it. I maintain that this latter property explains the widespread intuition that meaning is normative; when we make claims about what expressions mean we are commonly also intentionally conveying

normative information as well. This normativity has had a tendency to stick to meaning, even though meaning isn't itself normative.

One might wonder then why it hasn't also stuck to physics and arithmetic, since these subjects have the same degree of normative involvement as meaning. I suspect the reason why is that it is relatively easy to see that physics is descriptive because we are relatively clear about what it is supposed to describe; physics is about the physical world. Arithmetic too, has a subject matter which, at least at a superficial level, is obvious. Arithmetic is about numbers, operations on numbers, and the properties and relations that hold of them. Describe those correctly, and the arithmetical theory is true. Describe them as being other than they are and it is false. However, historical confusion regarding the aims of a theory of meaning has obscured meaning's transparently descriptive base.

1.3 The Normative Status of Semantic Rule-following

1.3.1 Semantic Rules as Constitutive Rules

If we take on board the claim that meaning is grounded in semantic rules, we must determine what kinds of rules semantic rules must be to play the role attributes to them. Drawing on Searle, we can answer this question by distinguishing between constitutive and regulative rules.²⁰ Constitutive rules are rules that create new kinds of actions. They create kinds of actions that could not be performed if the rule did not exist or were not in force. As a result, the activity metaphysically depends on the rules. This stands in contrast to regulative rules which merely regulate an activity whose existence is

²⁰ Searle (1963)

independent of the rules. Paradigmatic cases of constitutive rules are rules of a game. In order for my action to count as scoring a touchdown, I must carry the ball across the goal line. But that only makes sense against a back drop of the rules of football, since without the rules of football, there are no goal lines. Contrast this with the case of a no-taunting rule. There is nothing about understanding taunting as an act that requires appeal to any rule. To taunt just is to provoke someone with insulting language, and that can be done whether or not there is a rule in place sanctioning provocation by insult. So, a no-taunting rule is a regulative rule.

Semantic rules are supposed to play two primary roles. They make it the case that an expression has a meaning, and also make it the case which meaning it has.²¹ Thus, the semantic rules make meaningful language use metaphysically possible. In light of this, it is most natural to understand semantic rules as constitutive rules.

With this in hand, a natural picture emerges regarding how it is that the semantic rules give rise to semantic norms. The rules of chess, in bringing about the possibility of chess gameplay, give rise to various obligations, e.g., that you should checkmate the opponents king. Similarly, semantic rules, in making meaning possible, give rise to various obligations, e.g., that one ought to apply the expression “dog” to all and only dogs.²² In bringing about a certain practice, presumably constitutive rules bring with them prescriptions regarding

²¹ Glur (2009), Searle (1963)

²² This raises issues about how we ought to be formulating the relevant norms, and it is contentious the way to go on this. There is in the literature some who have expressed skepticism about whether this can be done, but I’m going to take up this question in a later section where I will offer novel reasons for being skeptical that we can articulate plausible semantic norms. Nothing I say here turns on the deontic modality used in my example. I use “ought” only for simplicity of comparison.

that practice. But before we can evaluate whether or not the semantic rules are normative in a way that is problematic for the naturalist, we must distinguish between the two kinds of norms regarding practice we could take to follow from the semantic rules.

- (1) Norms regarding engagement in the practice.
- (2) Norms regarding execution of the practice.

If we take our linguistic community to be competent speakers of English, then a norm of the first kind will be a norm regarding *whether or not* one ought to make use of English expressions in their communicative exchanges with other speakers. A norm of the second kind will be a norm regarding *how* one ought to make use of English expressions when communicating.

Norms of the first kind are clearly not the right kinds of norms. We should not be tempted to think that the normativist takes the constitutive semantic rules to give non-instrumental reasons to speak English. It is certainly clear that the rules of chess (our guiding paradigm for constitutive rules) entail nothing about whether or not one should play chess on their own. The rules give you a reason to play the game only perhaps insofar as you regard the rules as constituting a fun or challenging game and you want to play a fun or challenging game. Similarly, in the natural language case, the rules of English constitute an activity participation in which enables you to coordinate with others and among other things exploit them as sources of information about the world.

The conventional status of the semantic rules makes this possible, but they only give you reason to make use of the expressions they make meaningful

insofar as you find yourself wanting to coordinate with speakers of that language. What rationalizes my pointing to an object and applying the expression of a certain language? Well, if one is in the company of English speakers, they will conclude certain things about the object given the application of certain expressions. If the expression “gift” is applied, they will conclude that it is a present. However, if one is in the company of German speakers, they will conclude that it is poison. So, the reasons that I have to apply certain expressions of a language to objects will be reasons that arise from what I’m trying to get my community members to do. If I want them to stay away from the object and they are Germans, then I should apply the expression “gift.” So, instrumentally, the way I find out what I should do crucially depends on facts about what language conventions they happen to be following.

Norms of the second kind, however, seem more promising. It does seem like the rules of games like chess on their own give rise to obligations for players to play the game in a certain way. For example, it seems right to say that the rules of chess prescribe that one ought to put her opponents king in check, or that one ought to avoid stalemate. So, in characterizing what it is to be a player of chess, one must account for certain normative pressures in force for the player. The trouble, however, is that it is precisely here where the analogy with chess and related games breaks down. It is my aim in the next section to trace out carefully why this is so. Unfortunately for the normativist, semantic rules give rise to norms regarding how to use language only in the presence of other normative assumptions. Thus, semantic rules lack grade 1 normative involvement and the meta-semantic version of normativism fails to pose a serious challenge to naturalism about meaning.

1.3.2 Norms of Execution are Instrumental Norms

It will be useful to begin by reflecting a little more carefully about games and the role that the comparison to gameplay occupies for the normativist. Take for example the card game *Bang*. When playing *Bang*, each player is assigned one of four roles: sheriff, deputy, outlaw, and renegade. The sheriff and deputy win when the outlaws and renegade are eliminated from play. The outlaws win just as soon as the sheriff is eliminated, and the renegade wins when he is the last one in play. Interestingly the rule specifying the various roles make use of normative vocabulary:

Each player has his own goal: Sheriff: *must* eliminate all the Outlaws and the Renegade, to protect law and order. Outlaws: *must* kill the Sheriff, but they have no scruples about eliminating each other to gain rewards! Deputies: *must* help and protect the Sheriff, and share his same goal, at all costs! Renegade: he *must* be the last character in play.” (emphasis mine)

So, to be the sheriff is to be such that one ought to eliminate the outlaws and the renegade. This is significant. Something will not count as being a sheriff unless it is subject to certain norms. The rules of *Bang* institute irreducibly normative roles for the tokens governed by the rules. Pieces of the game, such as “bang” and “beer” cards are defined in relation to their role of eliminating or preventing the elimination of players with certain roles. Thus, they are cards that (say) sheriffs and outlaws ought to use in certain ways. *Bang* is not unique in this regard. Many games are such that their rules bring about irreducibly normative roles for the various pieces governed by the rules of the game. But not all games have this feature. Take for example, chess. In chess,

the pieces of the game, bishops, pawns, rooks, etc., have no irreducibly normative role. The piece movement rules are not defined in terms of how it is that they contribute to the realization of check-mate board configurations. Consider the movement rules for bishops and rooks²³:

1. The bishop may move to any square along a diagonal on which it stands
2. The rook may move to any square along the file or the rank on which it stands

However, the rules of the game still specify a goal or a victory state for each player:

The objective of each player is to place the opponent's king 'under attack' in such a way that the opponent has no legal move. The player who achieves this goal is said to have 'checkmated' the opponent's king and to have won the game.²⁴

Thus, to be a player of chess is to be such that one has as an objective to bring it about that his opponents king has no legal move available. To be a player of chess is thereby to be under a certain normative pressure to checkmate one's opponent. Games like *Bang* and Chess come apart in that one and not the other specifies irreducibly normative roles for the pieces used to play the game. However, they share the important feature that they have rule-specified goal or victory states. As a result, it is constitutive of what it is to be a player of the game that one is subject to certain norms.

²³ FIDE Laws of Chess Rule 3.2 and 3.3

²⁴ FIDE Laws of Chess Rule 1.2

There is a third class of games that lack even this feature: rule-specified victory states. Paradigmatic games of this sort are open-world video games. When working through the tutorial at the beginning of the game, the player is presented with rules merely for transforming game states. Pressing the A-key swipes a sword or shoots a gun. Pressing the direction pad moves the avatar in certain directions. Pressing the start-key opens up an equipment menu. But once the player is instructed in these game-state transformation rules, the player is left to pursue what missions the player wants. The game itself doesn't privilege any of the various available missions or quests. In these games, if what you want to do is lurk in the shadows and leap out and strike monsters, then the player can figure out how to use the rules to make that happen. If what he wants to do is rescue the princess in distress then the player can figure out how to exploit the rules to make that happen. He may be particularly satisfied when he does, but it is not as if he cannot play the game given that the game hasn't provided a goal for him to pursue. In games like these, it is not constitutive of what it is to be a player that one ought to do anything in the game. Whatever the player ought to do is crucially a function of the aims and interests the player brings to the game.

We may now ask what the status of the norms are in the various types of games. In the open world games, it seems the relevant norms will be instrumental. This is due to their being a function of the player's contingent aims and interests. In games like Bang and Chess, however, the kinds of norms specified in the rules are non-instrumental. If one is dealt the sheriff role, regardless of whether or not one wants to eliminate the outlaws, there is still a normative pressure to eliminate the outlaws and bring about the victory state. Similarly in chess, whether one wants to win or not, there is an obligation to

checkmate one's opponent. One simply is not a player without being subject to such a norm.²⁵

This point is worth emphasizing. In the case of games like *Bang* and chess, the non-instrumental status of the norms generated by the rules is a direct consequence of the role they play in realizing the rule-stipulated victory states. This is a crucial point. What does all the work in generating the non-instrumental normative consequences of the rule is the presence of a victory state; something that the player occupying that role *ought* to bring about.

This gives us a criterion by which to judge the plausibility of the meta-semantic version of normativism. If the meta-semantic normativist is right, we should expect to find victory states generated by semantic rules. The trouble is when we go looking, we won't find any. This prompts one of two responses. We can reject the meta-semantic normativity thesis or we can try to find a way for the semantic rules *simpliciter* to generate normative consequences without invoking victory states. In what follows, I try to cut off the second line of retreat from the outset. I will argue that the only way that constitutive rules *simpliciter* could generate norms is by invoking victory states. Then, by showing that semantic rules lack victory states, we will be forced to conclude that semantic rules are not intrinsically normative and thereby reject the meta-semantic normativity thesis. The upshot is that while semantic rules are like

²⁵ Here it is important to note that the kinds of obligations that the rules give rise to are *prima facie* obligations. Certainly, we want to accommodate the fact that individuals can engage in gameplay to lose, otherwise it would be impossible to make sense of say, parents playing games with their children. One can certainly find one's obligation to realize the victory conditions of the game overridden by other relevant considerations, like building the self-esteem of your opponent. But it does not follow from the obligation being overridden by other (perhaps worthy) desires of the agent that any aims or desires of the agent figure into that in virtue of which you have the reason to try and win. See e.g., Whiting (2007) for a discussion of this in greater detail.

games, they are importantly unlike games such as Bang and Chess. Rather language is more like an open-world game, where the semantic rules facilitate psychological state transformations allowing agents to realize their contingent communicative aims and interests.

In the interest of clarity and for ease of discussion, it will be helpful to have a more formal reconstruction of the argument before us:

1. Constitutive rules have normative consequences *simpliciter* only when they specify victory states.
2. Semantic rules do not specify victory states.
3. Therefore, Semantic rules *simpliciter* lack normative consequences.

1.3.3 Defending P1

How does one go about showing that P1 is true? The plan is to take something that is clearly instrumental and game like, a sort of proto-chess and add back small pieces one at a time so that it will be apparent where the problematic kind of normativity is coming in. The aim of the following thought experiment is to show that the stage at which this occurs is adding a rule-specified victory state.

Imagine that two people, call them Kim and Mike, are abducted by aliens. They are taken to a room in the alien space ship with a single table and 2 chairs. On the table is an 8x8 board where the tile colors alternate between black and white. Arranged on the back two rows of each side of the board are strange metal figures black and white, and each side of the board has the same pieces in the same configuration. Kim suspects it's a game and asks the aliens how to play. The aliens give the following instructions to Kim and Mike and leave.

Instructions: in Zorbs, each player takes turns. Your pieces are the metal figures on your side of the board. You can move exactly one piece per turn. When it is your turn, by touching one of your pieces, the board will illuminate and a well-defined move will be a movement of one of your pieces to one of the tiles on the board that has its matching color. When you move a piece to one of its correspondingly illuminated tiles your turn is over, and it becomes your opponent's turn. If one of your opponent's pieces occupies one of the illuminated tiles corresponding to your piece, you can move your piece to that tile and remove your opponent's piece.

This is a very boring and uninteresting game. Nevertheless, Kim and Mike start to play Zorbs. The only rules that they have been given by the Aliens are rules for transforming one game state into another. What kind of reason do the rules give say Kim to move one of his pieces to one of the correspondingly lit up tiles? It should be clearly instrumental. Kim has a reason to move his piece either in the interest of realizing board states he finds worth realizing (perhaps out of curiosity), or in the interest in coordinating with Mike such that Mike can realize board states he is interested in realizing.

First Addition: Kim and Mike start playing and exploring board states. As they play, they begin to make inductive generalizations about how the pieces move. For example, they notice after a while that one of the pieces has only ever had tiles diagonal to it illuminated and conclude that this piece moves diagonally. Suppose after a while, they form an inductive generalization of this kind for every piece on the board and the

generalizations that they make are extensionally equivalent to the rules of movement for chess pieces.

Does adding exhaustive piece movement rules change the instrumental normativity of the rules? No. At this stage as in the previous, they only facilitate the realizations of board states of interest to Kim and Mike. The only difference is that Kim and Mike can predict what tiles will illuminate when the various pieces are pressed.

Second Addition: as Kim and Mike keep playing, they end up in a board state where the board displays a cat emoji and resets to original position. They think “that’s an interesting state...never seen that one before.” After further play and more inductive generalizations they infer that cat-state is realized when a particular piece of Kim’s only has tiles illuminated that intersect with tiles that will illuminate for Mike’s pieces. Out of curiosity, they attempt to get into the same board state but this time with Mike’s piece instead of Kim’s. When they realize this board state, the board displays a dog-emoji and reverts back to original position. This is extensionally equivalent to checkmate in chess.

By adding states extensionally equivalent to checkmate have we recovered enough of chess to satisfy the normativist? Not yet. They could certainly become interested in realizing cat-state or dog-state, but this would be in the same way that they were interested in trying to bring about the board states when they were exploring board possibilities. At best we get that one of the things they like about the game is that there are things you can do to have the game display cat and dog emoji’s, and that they know how to use the rules

to bring that about. But even if we're not there yet, it looks like we're getting close.

Third Addition (Anti-coordination): Each time Kim plays he tries to realize dog-state and each time Mike plays he tries to realize cat-state.

We have already seen that there is no non-instrumental normativity when we have a state of the board that we mutually agree we would like to see if we can attain that board state through the rules. But the same can be said for anti-coordination. Lack of mutual agreement won't change things. Whether or not Kim should bring about dog-state depends on his having an interest in doing so. So, this won't get the normativist what he needs and perhaps we weren't expecting it to, but we need anti-coordination to be in place because it allows us to introduce:

Fourth Addition (New-comers): As Kim and Mike sit down to play Zorbs, the aliens walk in with two new abductees: Todd and Sara. The aliens forbid them from communicating with each other and instruct Kim and Mike to carry on as they've been. Kim tries to bring about dog-state and Mike tries to bring about cat-state. Over the course of the next few days, Todd and Sara make the following observations:

1. Kim and Mike are each respectively are trying to bring about dog-state and cat-state.
2. When Kim brings about dog-state she gets really happy, and when Mike brings about cat-state she gets upset and says things like "you got me this time but I'll get you next time!" (mutatis mutandis for Mike)

3. They make all the same inductive generalizations Kim and Mike made about piece movement and how to realize cat-state and dog-state.

One day the aliens come back and take Kim and Mike away. Todd and Sara begin to play. Todd brings about dog-state and celebrates saying “I can’t believe I won!” and Sara walks away upset. Now suppose that the aliens have the following exchange with Todd:

Aliens: “why did you get so happy when you realized dog-state?”

Sara: “because I won.”

Aliens: “What do you mean you won?”

Sara: (staring incredulously) “well...I’m supposed to use my pieces to bring about dog-state. I’m black and those are the rules, right?”

Now it seems we are back at full blown chess. We should expect that what’s missing for the normativist re-enters the picture and that’s what seems to be the case. Todd felt a normative pressure to bring about dog-state and was delighted when he satisfied what he took to be a demand the rules made of him. But what is different about Todd and Sara’s situation?

The difference seems to be crucially how they came to learn the game. When Kim and Mike learned Zorbs it was clear what the rules were, and there was nothing in the rules that told them anything about what states they should or shouldn’t bring about. Kim and Mike had a goal state only from a perspective that sat above the rules. But when Todd and Sara come to learn the game, they take the goal of realizing cat-state and dog-state to be supplied

by the rules themselves. As a result, they had internal rule-specified goal or victory states.

1.3.4 Defending P2: Semantic Rules Do Not Specify Victory States

In order to fix our judgments, it will be useful to look at the semantic rule-following constitutive of a formal language like classical propositional logic (hereafter PL) and then see whether our judgments about the formal language extend to a natural language like English. Applying the normativist's thesis to the case at hand, to claim that formal semantic rules on their own have normative consequences is to advance something like:

Where Γ is the set of semantic rules governing the meaning of the logical connectives and atomic sentences of PL, and p is some privileged theorem of PL:

Norms of Derivation: $\Gamma \models$ one ought to derive p

The trouble is, no instance of the above principle seems true. When handed the semantic rules for PL, if they are essentially normative in the way the normativist predicts, then it should make sense to ask the question, "what should I be trying to get on the bottom line of my derivation?" and expect some candidate substitution instance of p to be made salient by having someone direct the inquirer's attention back to the semantic rules. However, this simply cannot be done. The rules of PL don't accord a privileged status to any of its theorems. For example, the rules of PL don't entail anything special about the law of excluded middle. It's not as if it is the point of PL to try to produce a derivation of the law of excluded middle. This marks an important dissimilarity between the kinds of constitutive rules operative in a game like

chess and those operative in a formal language. It makes perfect sense to ask “what should I be trying to do with my bishop” and expect an answer from being pointed back to the clause in the rule book that talks about checkmate. It makes no sense to ask “what should I be trying to derive” and expect an answer to come from being pointed back to the clauses for negation and disjunction.

But we should have expected this given an understanding of the nature of the formal semantic rules. It does follow from the rules that I can move from certain sentences to other sentences and preserve truth. For example, given the semantic clauses, I know that I can move from $p \ \& \ q$ to p and preserve truth. So, the rules characterize certain kinds of sentential transformations; the truth-preserving ones. But we’ve already observed that rules that merely characterize state-transformations are insufficient to bring about the required strength of normativity. What we have in the case of the formal semantic rules is quite like what we had in the thought experiment at the stages **Introduction** and **First Addition**. Suppose in a game of Zorbs, one was given a board state with the equivalent of a bishop near the center tiles. All the tiles on the diagonal will be illuminated. This is analogous to being given an instance of modus ponens where the antecedent and consequent are complex and one has the antecedent. Deriving the consequent then is analogous to moving the piece. If I have $A \ \& \ B$ for complex A and B , I can either get A or I can get B , in the same way that if I had the Zorb equivalent to a pawn, I can move it up once or move it up twice.

Someone might hold that it is problematic to state the norm in terms of some particular privileged theorem, and that instead we should understand the rules to give pressure to get into all the tautological states. This would avoid

the worries I raise above, but it strikes me as implausible. There can't be normative pressure to prove all the theorems. Not only does that not square with the phenomenology of the relevant rule-comprehension (agents will not report that they feel such pressure when presented with the rules) but such a norm would be in principle unsatisfiable. There will always be more theorems to prove.

Now there are clear cases where the rules might give rise to normative pressure to derive certain theorems. For example, suppose that the rules were given to an intuitionist. Then, perhaps in trying to build a case against the classical logician he would need to use the rules to show what he considers to be absurd theorems you can use the rules to derive; e.g., $\sim\sim p \vDash p$. But it is clear that the rules are only instrumentally normative. Following them contributes to the realization of his aim of demonstrating the bad features of classical logic, in much the same way that the piece movement rules in Zorbs contribute to the players aims of realizing certain board states. This brings out what was missing by simply appealing to the semantic rules in answering the question "what should I be trying to derive?" We also need to appeal to the agent's derivational interests. If she wants to show why classical logic is absurd, then maybe she ought to use the rules to show $\sim\sim p \vDash p$. If she wants to show why the material conditional has certain paradoxical features then perhaps she should use the rules to show $p \vDash q \supset p$. Or if she wants to explore what the various theorems are, then she should use the rules to derive as many theorems as suits her. But none of these interests or goals can be given by the rules themselves. Thus, it seems the kinds of semantic rules constitutive of the meaning of the expressions in the formal language fail to be intrinsically normative.

I think very similar things should be said about semantic rules that constitute the meaning of natural language expressions. But it is much less straightforward in the natural language case as to why. While it is somewhat clear in the formal language that the rules themselves will not provide an adequate answer to the question “what should I be trying to do?” speakers (of English at least) are much more inclined to regard appeals to the semantic rules as constituting good and informative responses to such a question. So, in order to see whether or not we should extend our judgments from the formal language case to the natural language one, we need to bring into sharper relief just what exactly the semantic rules will look like.

Different conceptions of semantic theories will correspond to different conceptions of what the semantic rules look like. If you have an object based semantic theory where expressions are just assigned semantic values that are supposed to compositionally combine through functional application, then the semantic rules will articulate a relation between an expression and its model theoretic semantic value; for example, we might expect to see a rule like (as an instance):

Objectual CC: ($\llbracket \text{chair} \rrbracket = \{x: x \text{ is a chair}\} \rightarrow (\forall y (\text{“chair” correctly applies to } y \text{ iff } y \in x))$).

However, Objectual CC is clearly not a prescriptive rule. It does not tell a speaker to apply the expression “chair” to something just in case the object is in the class of objects comprising the extension of the expression. In fact, it doesn’t tell the speaker to do anything. As a result, it is difficult even to make sense of the idea that someone could “follow” a rule such as Objectual CC. And if it provides no directions then it certainly can’t be responsible for setting

a goal or directive for speakers of the language. So, if rules like Objectual CC are the rules that are meaning determining, then they won't generate the sort of normativity the normativist needs. This is significant because it puts pressure on the normativist to give up on certain conceptions of semantics. It is an interesting result if endorsing the normativity of meaning amounts to a rejection of objectual semantics. That much certainly wasn't obvious from the outset.

What conceptions of semantics are more amenable to the normativist line? Presumably the most simpatico conception is an inferential role semantics. In fact, many normativists have argued that the best versions of normativism require an inferential role semantics.²⁶ However, while it seems like a promising move to appeal to meaning as inferential role, ultimately, I think it won't provide much help. The inferential role semanticist maintains that meaning is exhausted by the web of inferential patterns associated with a sentence. In other words, meaning is exhausted by the sentences which you could infer from it and those from which you can infer it. This explains why derivational systems are considered classic exemplars of an inferential role semantics. One need not think of the derivational system as giving you a semantics of the logical expressions in the formal language, but those are the kinds of things an inferential role semanticist is going to point at and say that knowing those things is *all there is* to knowing the meaning of the logical expressions. For example, knowing

$$\frac{P \quad Q}{P \wedge Q} \quad (\wedge\text{-I}) \qquad \frac{P \wedge Q}{P} \quad (\wedge\text{-EL}) \qquad \frac{P \wedge Q}{Q} \quad (\wedge\text{-ER})$$

²⁶ See e.g., Peregrin (2012)

is all there is to knowing the meaning of conjunction. However, as we have already seen, one can use a derivation system without running afoul of the problematic type of normativity. So, we shouldn't expect taking the further step and identifying meaning of the expression with the inferential rules to help the normativist at all.

In light of this, we can extend our judgments about the normative status of the formal semantic rules to the case of natural language semantic rules. Just as formal semantic rules have normative force only in relations to the agent's derivational aims, so too do natural language semantic rules have normative force only in relation to agent's communicative aims. That's really all the semantic rules are: tools for transforming the psychological and behavioral states of others. But the target psychological and behavioral states—the states that I should try and get people into—are up to me. If I want to get people to think about Dummett and be understood as trying to get them so to think, then I should make use of his name in my speech acts to others. And there's no obstacle here to the project of giving a naturalistically respectable theory of meaning.

1.4 A Worry

I want to address a worry that the reader might have regarding the argument just made. The worry is that a crucial data point has been ignored, one that threatens to undermine the results from the previous section? Semantic rules, after all, are violable. They can be broken. And it is a characteristic feature of normative rules that they can be violated. Descriptive rules by their very nature cannot, so something must have gone wrong. I think the right response to the violability data is to say that the alleged violations are

not really violations at all. The data is a result of loose ways of characterizing what is going on. Suppose we have the following situation. A student is doing a homework problem for an introductory logic class and makes the following move:

- ⋮
- 3. $A \supset B$
- 4. $\sim B$
- 5. A (3,4) By, Modus Ponens
- ⋮

Here we clearly have someone making a modus tollens mistake, but it makes little sense to say that the logical law of modus tollens has been violated. What would count as a real violation of modus tollens would have to be a model where (3), (4), and (5) are all true. But if there were such a thing then modus ponens would not be a logical law. Given that there is not, then modus tollens is a law. We can say the same things about mathematical mistakes. Students in basic arithmetic courses make mistaken calculations all the time, for example, a child might divide 42 by 6 and get 7. But this doesn't amount to a violation of the laws of arithmetic. The best way to think about what is going on in these two cases is a lack of understanding about what the relevant laws are. When speakers misapply expressions, it is the best explanations of what is going on to say that they lack a certain understanding of what the meanings are. In fact, our ordinary practice of speech correction seems to presuppose this. If the misapplication wasn't a result of misunderstanding what the meaning is, why would we regard as an appropriate response statements like "I don't think that word means what you think it means."? So just as we shouldn't say in light of logical and arithmetical mistakes that the laws of arithmetic and logic are violable, so too should refrain from making the

inference in the semantic case. Thus, we can square the possibility of a semantic mistake without requiring what its being mistaken as problematically normative.

1.5 Preliminary Conclusion

I've argued in the previous sections that the most plausible way of understanding the normativist thesis is as a meta-semantic thesis. I then showed that on the normativist's preferred meta-semantic account, semantic rules fail to introduce the kind of normativity that could jeopardize the prospect of giving a naturalistic theory of meaning. Since semantic rules fail to specify goal or victory states they lack constitutive normativity. Thus, we should reject the claim that the normative implications of meaning facts impose a serious constraint on an acceptable theory of meaning. But we need not stop here. I think it can be shown that there is in fact no special normative relation between what an expression means and how it ought to be used. This is the topic of the next section.

1.6 Articulating the Meaning Norms

Ordinarily when someone makes a claim to the effect that x is normative, it is appropriate to expect an answer to the question "What is the norm?" If I claimed that the good is normative, there would be a straightforward way of saying what the norm is: one ought to promote the good. Inability to provide a satisfactory answer will be a serious consideration against the claim, if not an outright refutation of it. In the foregoing, I've been acting as though it is unproblematic to remain neutral on what the relevant semantic norms might

be. Whatever the normativist had in mind, it is doubtful that any non-instrumental norms would be forthcoming for the reasons I gave. But regardless of the status of the norms, it is assumed that some plausible norms will be forthcoming.

There is quite a bit of controversy as to how exactly we should formulate the norms. Some have argued that we should understand meaning norms as claims about how we “ought” to use expressions.²⁷ Other have argued that the norms are much weaker, and instead we should understand them in terms of permissions to use expressions in certain ways.²⁸ The dialectic has proceeded largely case by case. Someone offers up a candidate norm, and those that demur offer up counterexamples. However, if real progress is to be made in answering the question, we need a principled and systematic way to evaluate the range of possible norms. I want to close by offering a systematic way of examining the logical space as well as some novel considerations in favor of skepticism about meaning norms. This will underscore the results from the previous sections. Just as there is no special connection between logical consequence and how we should manage our beliefs, there is no special connection between what an expression means and how we should use that expression.

1.6.1 Logical Space

In order to proceed in as principled and systematic way as possible, we need a way of identifying all of the candidate forms the meaning norms might take. Thankfully, a good deal of work in mapping out the logical space can be

²⁷ Hattiangadi (2006)

²⁸ Whiting (2007)

done by making appropriate adjustments to some of the work John McFarlane has done within the context of the normativity of logic.²⁹ In his unpublished manuscript, McFarlane provides a taxonomy of bridge principles. A bridge principle takes the form of a material conditional, the antecedent of which states, for McFarlane facts about logical consequence, and the consequent contains a normative claim concerning the agent’s doxastic attitudes whose contents stand in the said logical relations. It will be useful for our purposes to understand bridge principles in a similar fashion, only we will take the antecedent to state facts about an expression’s meaning and the consequent to contain a normative claim about speakers’ use of the expression. So, we can take the relevant bridge principles as instances of the following schema. Where e is an expression and F gives the meaning of that expression:

$$\forall e (\text{if } e \text{ means } F) \rightarrow \phi$$

Taking this schema as a point of departure, we can follow McFarlane in introducing three parameters. Each parameter allows for multiple discrete settings. The range of possible combinations among these parameters exhausts the logical space of bridge principles. The possible parameters are as follows:

1. Bridge principles may differ in the deontic operator they deploy: Does the normative constraint take the form of an ought (o), a permission (p) or merely of having (defeasible) reasons (r)?
2. What is the polarity of the normative claim? Is it a *positive* obligation, permission, or reason to use a certain expression given its meaning, (+)?

²⁹ McFarlane (2004)

Or rather is it a *negative* obligation/permission/reason *not to use* it a certain way given its meaning, (-)?

3. Different bridge principles result from giving the deontic operator different scope. Let O stand generically for one of the above deontic operators. Given that the consequent of a bridge principle will typically itself take the form of a conditional, the operator can take
 - Narrow Scope with respect to the consequent (N): $(P \rightarrow O(Q))$
 - Wide scope (W): $O(P \rightarrow Q)$
 - It can govern both the antecedent and the consequent of the conditional (B):
 $O(P) \rightarrow O(Q)$.

These settings can be combined so as to generate a total of eighteen bridge principles. The symbols in parentheses associated with each parameter, combine to determine a unique label for each of the principles: The first letter indicates the scope of the deontic operator (N, W or B), the second letter indicates the type of deontic operator (o(bligation), p(ermissions), r(easons)), and “+” and “-” indicate the polarity. It will be useful to include a table setting out all of the possibilities. Table 1 does just this and is an adaptation from McFarlane’s.³⁰

³⁰ McFarlane (2004, 7)

Table 1: e means F, then $\forall x \dots$

N Deontic operator embedded in consequent.

o Deontic operator is strict obligation (ought).

No+ if x is f , then you ought to apply e to x .

No- if x is not f , you ought not apply e to x .

p Deontic operator is permission (may).

Np+ if x is f , then you may apply e to x .

Np- if x is not f , you are permitted not to apply e to x .

r Deontic operator is “has (defeasible) reason for.”

Nr+ if x is f , then you have a (defeasible) reason to apply e to x .

Nr- if x is not f , you have a (defeasible) reason not to apply e to x .

W Deontic operator scopes over whole conditional.

o Deontic operator is strict obligation (ought).

Wo+ you ought to see to it that if x is f , then you apply e to x .

Wo- you ought to see to it that if x is not f , then you don't apply e to x .

p Deontic operator is permission (may).

Wp+ you may see to it that if x is f , you apply e to x .

Wp- you may see to it that if x is not f , you do not apply e to x .

r Deontic operator is “has (defeasible) reason for.”

Wr+ you have reason to see to it that if x is f , you apply e to x .

Wr- you have reason to see to it that if x is not f , you do not apply e to x .

One notable difference between the bridge principles connecting meaning and norms of expression application and bridge principles connecting logical consequence and norms of reasoning is that there aren't any bridge principles of type B. This is because the antecedent of the embedded conditional in the linguistic case contains no mention of a speaker. As such, it would not be well formed to have a deontic operator governing just the antecedent of the conditional. This allows us to restrict the range of plausible bridge principles from eighteen to twelve at the outset.

We should also observe that in many cases the wide and narrow scope principles will be equivalent to each other. The only situations in which the wide scope and the narrow scope principles will diverge is when one is in a position to bring about whether an object is such that the application of an expression to it would be correct or not. For example, If I were in a position to construct a boat out of wood then I could see to it that the composition of the wood, nails, etc is or is not a boat, and thereby make it fall within or without the extension of that predicate. However, the artifact cases are the exception. I can't make it the case that a star is a star, or that a horse is no longer a horse. Given this connection between the wide and narrow scope principles, we can evaluate them more economically.

Some last comments regarding the table. First, the intended reading of "you ought not ϕ " is not "it is not the case that you ought to ϕ ," but instead "you are forbidden to ϕ ," which can be understood as "it is obligatory that you refrain from ϕ ing," that is, "it is obligatory that you see to it that you do not ϕ ."³¹ Second, the difference between obligations and reasons is that

³¹ McFarlane (2004, 9)

obligations are strict, whereas reasons are defeasible and *pro tanto*.³² McFarlane gives the following gloss on the difference:

If one ought to do A and does not do A, one is as a result subject to criticism; this is not so if one has reason to do A and does not do A. If one has reason to do A and no reason not to do A, then presumably one ought to do A. But often one has reason both to do A and not to do A. In that case, the stronger reason determines what one ought to do. One ought to go to class even though one has reason to take a walk instead, for one has an even more compelling reason to go to class. Conflicting reasons of this kind are the norm in practical reasoning; conflicting obligations, if they are possible at all, are much rarer. In a genuine case of conflicting obligations, one cannot emerge blameless. No matter what one does, one is remiss for failing to do what one ought to.

We will follow McFarlane in leaving it open whether it is possible for obligations to conflict in this way.³³

What considerations tell for and against the various possible ways of filling in the bridge principle? Here again I think it is helpful to look at the state of the dialectic in the case of the normativity of logic. Harman famously raises a number of considerations all aimed at ruling out any of the logical bridge principles. The considerations he raises have very close linguistic analogs which put similar pressures on ruling out many of the meaning bridge principles. So, first we will represent Harman's famous worries, and then we

³² Broome (2009, 79–81)

³³ McFarlane (2004, 9)

will identify the linguistic analogs. Afterwards, we will examine which bridge principles can be ruled out on the basis of these concerns.

1.6.2 Harman's Worries and Their Linguistic Counterparts

When trying to account for the normative connection between logical consequence relations and an agent's doxastic states, the following is a very natural first attempt. We reflect on the aims of theoretical reasoning and realize that theoretical reasoning aims to provide an accurate representation of the world. We accurately represent the world by having true, perhaps even knowledge-constituting doxastic attitudes. Our doxastic attitudes have contents and these contents stand in certain logical relations with one another. Having an awareness of these logical relations would appear to be conducive to the end of having true beliefs and so is relevant to theoretical reasoning. In particular, the logical notions of entailment and consistency seem to be important to the theoretical project. If I believe truly, the truth of my belief will carry over to its logical consequences. Conversely, if my belief entails a falsehood it cannot be true. Similarly, if what I believe (in general or in a particular domain) is inconsistent, they cannot possibly afford an accurate representation of the world; at least one of my beliefs must be false.

Thus, we might posit the following two principles as accounting for the relevant normative relation:

Logical Implication: One ought to believe the logical consequences of one's belief

Consistency: One ought to have consistent beliefs.

Harman gives a number of arguments against these naïve principles. However, we will only be concerned with two arguments he gives against Logical Implication.

The first argument he gives is as follows. According to Logical Implication there is something wrong with one's beliefs if there is a proposition logically implied by them that one does not already believe. In such a case one should either add the belief or give up one of the implying beliefs. But this cannot be right. Many trivial things are implied by one's beliefs which "it would be worse than pointless to add to what one believes."³⁴ If one believes p , that trivially implies $(p \ \& \ (q \vee \sim q))$, $(p \ \& \ (r \vee \sim r))$, $(p \vee ((q \ \& \ (r \vee \sim r)))$, etc. There is no point in "cluttering one's mind with all these propositions. This motivates the following meta-principle:

Clutter Avoidance: One should not clutter one's mind with trivialities.

The second argument he gives is as follows. Suppose one believes the foundational axioms of number theory. It is an open question in number theory whether Goldbach's conjecture is true. Suppose for the sake of argument that it is in fact true. It will be the case that it is a logical consequence of the axioms of number theory. And by Logical Implication, one ought to believe that Goldbach's conjecture is true. But he seems completely wrongheaded. One would have no reason to believe that Goldbach's conjecture is true given that it is an open question in number theory. This motivates another meta-principle:

³⁴ Harman (1986, 12)

No Remote Consequences: one is not obligated to form beliefs about remote logical consequences on one's beliefs.

Of these two meta-principles, No Remote Consequences has been the most difficult to satisfy. Most, if not all attempts to articulate plausible logical bridge principles have failed this constraint. So much for the normativity of logic. With this in place we can articulate analogous meta-principles constraining the range of plausible norms governing expression application. What is the linguistic counterpart of clutter avoidance? I propose the following:

Linguistic Clutter Avoidance: one cannot be under normative pressure to clutter one's speech with trivial or redundant applications of an expression.

Why accept this principle? Ordinarily, throughout one's day, one comes into contact with a wide variety of objects. At home, one sees the couch in the living room. At the office, one sits in a chair at a desk, and perhaps at the library one sees a number of books on the shelf. But it is an unreasonable constraint on a speaker that she be under normative pressure to apply the expressions "chair", "couch", "desk" and "book" to these objects. Supposing that I was in my office alone, I would feel no normative pressure to break my work silence to say of the various objects in my office that they are a chair or desk and so on. Any norm that requires this and similar patterns of linguistic behavior of a speaker should be rejected.

We now turn to the linguistic analog of No Remote Consequences. I propose the following:

No Remote Applications: one cannot be under normative pressure to make remote expression applications.

Let me say more about what I have in mind here. Some objects are such that I can't reasonably be expected to correctly apply an expression to them. This might be a result of there being an expression that would correctly apply but which is too complex. There are many complex predicates which can be constructed out of simpler ones which may be appropriate to apply on a given occasion, but given the complexity, it will outstrip my ability as a speaker to make use of it on such an occasion. In this case the expression is too remote and thus I can't be expected to correctly apply it. It might also be the case that I am aware of an expression that correctly applies to object, but the object is too remote making me unable apply the expression to it. Any norm that requires this and similar patterns of linguistic behavior of a speaker should be rejected.

With these meta-principles in hand we can start eliminating candidate bridge principles. No Remote Applications allows us undermine one third of the outstanding options:

Wr+: you have a reason to see to it that if x is f, you apply e to x.

Wr-: you have a reason to see to it that if x is not f, you don't apply e to x

Nr+: if x is f, then you have a (defeasible) reason to apply e to x.

Nr-: if x is not f, you have a (defeasible) reason not to apply e to x.

There are a number of complex English expressions that competent speakers of the language could in principle construct. But given their

complexity, they will be too remote for the speaker. Yet it is true of those expressions that they have the meaning that they do. Speakers will have no reason whatsoever to apply those expressions or refrain from applying those expressions in accordance with their meaning. Since it won't usually be in the speaker's power to make it the case that x is or is not f , this should motivate us to reject both Nr and Wr. Linguistic Clutter Avoidance pressures us to give up another two:

Wo+: you ought to see to it that if x is f , then you apply e to x .

No+: if x is f , then you ought to apply e to x .

I'm under no obligation to see to it that given the fact that "chair" means chair, that I apply the expression "chair" to the object I'm currently sitting on. Sitting in silence surely can't be a paradigmatic case of violating meaning norms. Moreover, if I were in a classroom with other competent speakers, it would direct me to draw public attention to mundane and trivial facts like there is a chair, there is a desk, there is a door, etc. If anything, that is something that we ought not do.

Now we have eliminated half of the contenders. What can be said for or against the rest of the available bridge principles? The remaining bridge principles with the strong deontic modal can be eliminated on the grounds that there are cases clearly where a speaker can apply an expression not in accordance with its meaning. A pedestrian case is sarcastic speech. If after a losing the championship game due to bad refereeing, I remark to my friends "what a great ref!" then I will not have applied the expression correctly. The referee certainly is not in the extension of good, but it won't be the case that I have done something I have a strict obligation not to do. And since I would

not be in a position to bring it about that the ref is not a ref, we have strong reason to give up:

No-: if x is not f , you ought not apply e to x .

Wo-: you ought to see to it that if x is not f , then you don't apply e to x .

We also have strong reason to give up both variants of the positive permissions:

Np+: if x is f , then you may apply e to x .

Wp+: you may see to it that if x is f , you apply e to x .

These two norms look quite plausible. After all, wouldn't it be odd to say that you can't apply an expression in accordance with its meaning? For most expressions this is certainly true, however there is a non-trivial class of expressions for which it is impermissible to apply them because they mean what they do. Paradigm cases include slurs and other pejorative expressions. As a result, we should reject these outright permissions. This leaves the two negative permissions:

Np-: if x is not f , you are permitted not to apply e to what is not f .

Wp-: you may see to it that if x isn't f , you don't apply e to what's not f .

In order for these principles to hold, the following kind of situation must not be possible: one where the speaker is obligated to continue a conversation in such a way that the only conversational recourse the speaker has is to apply an expression incorrectly. Most if not all of us have been in conversations we really couldn't get out of. It's the second constraint that seems tough to satisfy.

When are we ever in a conversational situation where we have to keep the conversation going in such a way that our application of an expression forces our utterances to be false? As hard to imagine as that may be, I think we can imagine such a case. Here I think two kinds of cases present themselves naturally: cases where a false presupposition has been accepted into the common ground and cases where one makes a false referential use of a definite description. While the second will seem fairly artificial, the first, I think, is fairly pedestrian. I think most if not everyone has found themselves in such a conversational pickle.

False Presupposition: Suppose that at a dinner party a conversation develops about the party guests. One of the conversational participants, Ted makes a remark about the woman Jones attended the party with. He says, “I just found out that Jones’s wife is a doctor.” As it turns out, the woman that Jones attended the party with is not in fact his wife although it would be very easy for someone to make that confused judgment. In order to avoid embarrassing Ted and to be cooperative, all the conversational participants fail to point out to Ted that she is not in fact Jones’s wife. There’s really nothing at stake so they let it slide. At a later stage in the conversation, Mary, one of the conversational participants who knows Jones is not married, needs make something known to the group about Jones’s partner. Given that it is in the common ground that she is Jones’s wife and that she is merely a partner isn’t, she really should use “Jones’s wife” to refer to her.

False Referential Use: Suppose two dignitaries—one foreign, one domestic—are having a conversation about international intrigue. The

foreign dignitary discloses that an international cabal has planned a devastating attack on domestic soil and the head of the cabal is a mole in the domestic nation's spy agency. However, time is of the essence. The attack is imminent and they must identify the spy as quickly as possible. There is in attendance a government official who knows the names and whereabouts of all the agency's active spies, and the foreign dignitary says, "Tell me who he is right now." The domestic dignitary scans the room, and what is most observationally salient about the government official is that he is drinking a martini. Moreover, he is the only individual in the room drinking a martini. So, he replies, "He is the man over there drinking a martini." Unfortunately for the domestic dignitary, he doesn't know that in fact the man's martini glass contains no alcohol; only water. So, his referential use of the expression was false, yet uniquely referential.

Now is it the case that he really had to use *that* definite description? Weren't there plenty of equally good and true descriptions in the vicinity that would have done just as well. Given the observational salience of the martini drinking, the high stakes of the situation, and the fact that it allowed him to pick him out uniquely, it seems plausible to maintain that it would not have occurred to the domestic dignitary to say anything differently. Or if something else might have, we could just as easily build it into the case that the only properties whereby the government official could be uniquely distinguished are by his imbibing the contents of the martini glass. And since the domestic dignitary has no evidence that the contents are not alcoholic, why think it would have occurred to him to think he should say anything other than he did?

I find these considerations persuasive, but I recognize that they aren't decisive. However, I think even if these cases fail to rule out the weak outright negative permissions it is worth remarking that it is an interesting result that norms these weak are the most plausible norms forthcoming from the meaning facts. And if the only norms forthcoming from meaning facts are permissions to say nothing at all, then it's a wonder why anyone ever thought that meaning had a special relation to communication. At best, it is instructive for when I am permitted *not* to speak. And, if only instructive for when I can refrain from speaking, why should the normativity of meaning have ever been thought to introduce a substantive constraint on what counts as an acceptable theory of meaning.

But I take these considerations to at least draw out a very close connection that between the normativity of logic and the normativity of meaning. What on the surface looks like an obvious and straightforward normative relationship, becomes quite slippery if not impossible to make clear. In light of the results from the previous section, this should come as no surprise. The meta-semantic facts simply aren't normative facts. Any norms governing the use of our expressions then simply won't have a semantic or meta-semantic origin.

1.7 Conclusion

While “Meaning is normative” has become a popular dictum in the philosophy of language, and some proponents of the dictum have appealed to the normativity of meaning in order to demonstrate the hopelessness of a reduction of semantic to non-semantic vocabulary, we have seen that on the most plausible construal of the thesis, it fails to support such a strong conclusion. Moreover, we have seen that while initially plausible, there is no special relevance of an expression’s meaning to how it ought to be used. So not only should we reject the claim that meaning is normative in a way problematic for naturalism, we should be skeptical of the claim that meaning is normative altogether.

Works Cited

- [1] Bilgrami, Akeel. "Norms and Meaning." *Reflecting Davidson* (1993): 121–144.
- [2] Boghossian, Paul A. "The rule-following considerations." *Mind* 98.392 (1989): 507–549.
- [3] Blackburn, S. "The Individual Strikes Back," *Synthese*, 58: 281–30. 1984.
- [4] Bloor, David. *Wittgenstein, Rules and Institutions*. Routledge. (1997)
- [5] Brandom, Robert. *Making it Explicit*, Cambridge MA: Harvard University Press, 1994
- [6] ———. *Articulating reasons*. Harvard University Press, 2009
- [7] Bridges, Jason. "Rule-following skepticism, properly so called." *Varieties of Skepticism*. Berlin: de Gruyter (2014): 249–88.
- [8] Buleandra, Andrei. "Normativity and correctness: A reply to Hattiangadi." *Acta Analytica* 23.2 (2008): 177–186.
- [9] Davies, Martin. *Meaning, quantification, necessity: Themes in philosophical logic*. (1981).
- [10] Gamut, L. T. F. *Logic, Language, and Meaning, volume 2: Introduction to Logic*. University of Chicago Press, 1991.
- [11] Gampel, Eric H. "The normativity of meaning." *Philosophical Studies* 86.3 (1997): 221–242.
- [12] Glock, H.J. *A Wittgenstein Dictionary*, Oxford: Blackwell. 1996.

- [13]Glüer, Kathrin, and Åsa Wikforss. "Reasons for belief and normativity."
(2009).
- [14]———. "The normativity of meaning and content." (2009).
- [15]———. "The Normativity of Meaning and Content", *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.), URL =
<<https://plato.stanford.edu/archives/spr2018/entries/meaning-normativity/>>.
- [16]Goldman, A. 1998. "Strong and Weak Justification", *Philosophical Perspectives 2, Epistemology*: 51–70
- [17]Harman. G. *Thought*. Princeton University Press, Princeton, 1973.
- [18]———. Logic and reasoning. *Synthese*, 60:107–127, 1984.
- [19]———. *Change in view: Principles of reasoning*. M.I.T. Press, Cambridge, 1986.
- [20]Hattiangadi, Anandi. "Is meaning normative?" *Mind & Language* 21.2 (2006):
220–240.
- [21]———. *Oughts and thoughts: Rule-following and the normativity of content*. Oxford University Press on Demand, 2007.
- [22]———. "Some more thoughts on semantic oughts: A reply to Daniel Whiting." *Analysis* 69.1 (2008): 54–63.
- [23]Horwich, Paul. *Reflections on Meaning*. Oxford University Press, 2005.
- [24]MacFarlane.J. In what sense (if any) is logic normative for thought?
Unpublished, 2004.

- [25]Kaluziński, Bartosz. "Assessment, Scorekeeping and the Normativity of Meaning: a Reply to Kiesselbach." *Acta Analytica* 31.1 (2016): 107–115.
- [26]Kamp, Hans, and Uwe Reyle. *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and discourse representation theory*. Vol. 42. Springer Science & Business Media, 2013.
- [27]Kiesselbach, Matthias. "The normativity of meaning: from constitutive norms to prescriptions." *Acta Analytica* 29.4 (2014): 427–440.
- [28]Kripke, Saul A. *Wittgenstein on rules and private language: An elementary exposition*. Harvard University Press, 1982.
- [29]Lewis, David (1975). "Languages and language". In Keith Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*. University of Minnesota Press: 3–35.
- [30]McDowell, J. "Wittgenstein on Following a Rule," *Synthese*, 58: 325–363. 1984.
- [31]Peregrin, Jaroslav. "Inferentialism and the Normativity of Meaning." *Philosophia* 40.1 (2012): 75–97.
- [32][Rulebook] 2008 *Bang* (boardgame) Fourth Edition, daVinci games.
- [33][Rulebook] *Chess Handbook*, www.fide.com/fide/handbook.html.
- [34]Russell, Gillian (forthcoming). Logic isn't normative. *Inquiry*: 1-18.
- [35]Steinberger, Florian. "Three ways in which logic might be normative." *Unpublished draft of January 27* (2015): 2015.

- [36]Verheggen, Claudine. "Semantic normativity and naturalism." *Logique et Analyse* 54.216 (2011): 553–567.
- [37]Wedgwood. R. "The meaning of 'ought'." In R. Shafer-Landau, editor, *Oxford studies in metaethics*, volume 1, pages 127–160. 2006.
- [38]Whiting, Daniel. "The normativity of meaning defended." *Analysis* 67.294 (2007): 133–140.
- [39]———. "Is meaning fraught with ought?." *Pacific Philosophical Quarterly* 90.4 (2009): 535–555.
- [40]———. "What is the Normativity of Meaning?." *Inquiry* 59.3 (2016): 219–238.
- [41]Wikforss, Åsa Maria. "Semantic normativity." *Philosophical Studies* 102.2 (2001): 203–226.