Received: 10 June 2021 | Revised: 11 November 2021 | Accepted: 16 December 2021

DOI: 10.1111/phib.12262

ORIGINAL ARTICLE

ANALYTIC PHILOSOPHY WILEY

The structure of semantic norms

Jeffrey Kaplan

Department of Philosophy, University of North Carolina, Greensboro, North Carolina, USA

Correspondence
Jeffrey Kaplan, Department of Philosophy, University of North Carolina, 239 Curry Building, Greensboro, NC 27402-6170, USA.
Email: jikaplan@uncg.edu

Abstract

The normativity of meaning—introduced by Kripke in 1982, and the subject of active debate since the early 1990s—has been exclusively understood in terms of duty-imposing norms. But there are norms of another type, well-known within the philosophy of law: authority-conferring norms. Philosophers thinking and writing about the normativity of meaning—normativists, anti-normativists, and even Kripke himself—seem to have failed to consider the possibility that semantic norms are authority-conferring. I argue that semantic norms should be understood as having an authority-conferring structure, and show how this allows normativism about meaning to escape the two most popular arguments against it.

1 | INTRODUCTION

Semantic norms are norms that follow directly and invariably from meaning facts.[1] Are there such things? This is one way of posing the question at the heart of the normativity of meaning debate. If semantic norms exist, then normativity is an intrinsic feature of meaning that constitutes a constraint on theories of meaning. And that constraint, it has been thought, rules out behaviorist, dispositional, causal, and informational theories of meaning.[2] This is why the normativity of meaning matters: if meaning is normative, then all or some such theories of meaning are doomed. At first, the normativity of meaning was seen as obvious—nearly uncontestable.[3]

[1]This is a claim only about entailment, not about explanatory priority. From the fact that meaning facts entail norms, no conclusion can be drawn about whether the meaning facts or the norms are more fundamental.

[2]Wikforss (2001) uses "pure-use" as a label for theories of meaning that utilize purely descriptive explanatory resources, accounting for meaning in terms of how a term is used or how it is disposed to be used.

[3]Kripke 1982; Blackburn 1984; Boghossian 1989.

But in the decades since Kripke's initial presentation of it, the normativity of meaning has been challenged.[4]

The aim of this paper is to suggest—tentatively, at least—that semantic norms, and thereby the topic of the normativity of meaning as a whole, have been fundamentally misconceived, by both normativists and anti-normativists. The suggestion is that we should reconceive how semantic norms are *structured*. The point that I make in this paper is about the *form* of semantic norms. But this is not a trivial technicality. The reconceived version of the normativity of meaning that I introduce here (a) departs significantly from all versions mentioned in the philosophical literature by normativists, anti-normativists, and even Kripke himself, (b) avoids the central arguments that have been presented against the normativity of meaning, but (c) still constitutes a genuine form of meaning normativism that, if true, places a significant constraint on theories of meaning. If this paper succeeds, though, it will merely have pointed the normativity of meaning debate in a new direction. The fact that the version of normativism introduced here avoids the most popular and pressing objections to normativism as it has been previously understood surely counts in favor of the thesis that meaning is normative. But I do not offer any additional arguments in favor of the normativity of meaning. The purpose of this paper is to introduce a previously ignored version of the idea that meaning is normative, and not to demonstrate that meaning actually is normative.

2 | WHAT ARE THE SEMANTIC NORMS?

Normativists about meaning occasionally disagree about the order of explanatory priority—some holding that meaning facts should be explained in terms of semantic norms, and others holding that semantic norms should be explained in terms of meaning.[5] In this paper, I remain neutral on this issue by understanding normativism, roughly, as the view that meaning facts directly entail normative facts (this entailment relation being neutral as to direction of explanatory priority).

But what exactly are these norms that follow directly, without substantive auxiliary premises, from meaning facts? Here is a first attempt.[6] It follows from the fact that "cat" means *cat* that:

(A) "cat" ought to be applied to all and only cats.[7]

[4]Bilgrami 1993; Papineau 1999; Wikforss 2001; Glüer and Wikforss 2009; Glüer and Pagin 1998; Hattiangadi 2006, 2007, 2009. Those defending normativism include Rosen 2001; Whiting 2007, 2009, 2016; Ginsborg 2012.

[5]For overview discussions, see Glüer and Wikforss 2009, Glüer et al., 2022.

[6]In the following discussion, I will be following one trend in the philosophical literature, which formulates semantic norms in terms of *truth*. For instance, the first norm considered represents the fact that "cat" is *truly* applied to cats. But there is another historically prominent way of thinking about the normativity of meaning: in terms of *warrant*. Most prominent in this camp are those who advocate one form or another of inferential role semantics, such as Gibbard 2012, 14, 27, 204; Brandom 1994, 9, 626. I will put these views aside in this paper, though some of the problems outlined in the following section do apply to them. For an excellent discussion of the relation between truth-theoretic and warrant-theoretic conceptions of the normativity of meaning see Ginsborg 2018, 1002–1009.

[7]Hohwy 2001, 9–10.

As many have noted, this is too strong.[8] We are not required to apply "cat" to every cat we see, let alone to *all* cats. So consider instead:

(B) "cat" ought to be applied only to cats.[9]

Though more plausible on its face, this norm merely says how things "ought to be." As Hattiangadi (2009, 61) points out, this kind of "ought" can fail to be "agent-implicating."[10] By analogy, a knife ought to be used for cutting, but this does not entail that any agent ought to cut anything.[11] Here, is a third attempt, now involving a speaker. It follows from the fact that "cat" means *cat* in the language of a speaker *S* that:

(C) *S* ought to apply "cat" only to cats.

This norm avoids the problems that plagued (A) and (B), but it faces the following difficulty. What exactly is it for someone to "apply" a term? When one merely expresses a proposition—e.g., by asking the question of whether that proposition is true, or by asserting a conditional with that proposition in the antecedent—that predicates "cat" of something, one has, in some sense, applied "cat" to that thing. Understood this way, (C) is far too strong.[12] There is no prohibition against expressing false propositions in questions or the antecedents of conditionals. To avoid this, we can understand "apply" more restrictively: one has applied "cat" to something only when one *asserts* of that thing that it is a cat. But the norm should then be restructured to make this explicit. So, we arrive at the following: It follows from the fact that, in the language of a speaker *S*, asserting "That's a cat." of something means that that thing is a cat that:

(D) *S* ought to assert "That's a cat." only of cats.[13]

This looks very much like a truth norm for assertion. But it is worth emphasizing that the normativist thesis is not merely that *there are* norms broadly of this kind governing language use.

[8]Boghossian 2003; Bykvist and Hattiangadi 2007; Hattiangadi 2006, 2007, 2009, 58; Whiting 2009, 544.

[9]Whiting (2007, 137) holds something like this, though he presents it not as a semantic norm, but as a formulation of the normativism thesis itself: *w* means $F \rightarrow \forall x$ (*w* ought to be applied to $x \rightarrow x$ is *f*). Hattiangadi (2009, 61) points out a scope ambiguity, resolved by reformulating normativism as follows: *w* means $F \rightarrow \forall x$ (*w* ought to (be applied to *x*) $\rightarrow x$ is *f*).

[10]Humberstone 1971.

[11]It may be wondered whether this is really a problem, but what matters here is merely that it is taken to be a problem in the literature.

[12]Speaks 2009.

[13]Alternatively, we might wish to phrase this norm negatively:

(D')S ought not assert "That's a cat." of non-cats.

Nothing in what follows hinges on the difference between (D) and (D'). Indeed, none of the norms mentioned so far are phrased precisely as they are in the literature. (In particular, Hattiangadi 2009, 61 and Whiting 2009, 544 debate these issues not in terms of the semantic norms, but of the normativity of meaning thesis itself. Some further discussion of the differences between their formulations and the ones discussed here appears below.) This paper, however, presents a critique of the structure of semantic norms that is indifferent to the differences among all of these norms, including those mentioned here and those stated elsewhere in the normativity of meaning literature.

Rather, normativism about meaning can be expressed as the claim that the existence of a norm like (D) follows *directly*, *without any substantive auxiliary premises*, from the relevant meaning facts. Normativists hold not merely that language use is governed by norms—as virtually all human activities are—but also that such norms are intrinsic to meaning itself, and are not merely a reflection of the nature of particular speech acts, such as assertion. Though it is phrased in various different ways, most normativists take themselves to be defending and most anti-normativists take themselves to be attacking the claim that a norm like (D) follows directly from the fact that asserting "That's a cat." of something means that that thing is a cat.

3 | OBJECTIONS TO NORMATIVISM

If meaning facts directly entail norms like (D), then normativity is an intrinsic or essential feature of meaning that theories of meaning must explain. But there are two oft-repeated objections to meaning normativism understood this way.[14]

3.1 | The permissibility of lying objection

The first objection stems from the fact that, as some philosophers would put it, "we do not have any categorical semantic obligations: there are, for example, many circumstances in which it is permissible to lie."[15] Relevantly, that is, there may be no norms like (D). And if no such norms exist, then it is not true that the existence of such norms is directly entailed by meaning facts.

By itself, the occasional permissibility of lying does not show that there are no categorical norms like (D), but only that any such norms are *defeasible*, allowing for circumstances in which they are outweighed by other, non-semantic norms or considerations.[16] But the problem persists. It seems to some philosophers that there is not always even a *pro tanto* reason against asserting a falsehood. Rosen (2001, 621) offers the following example:

> You're playing math games with a child, but she's getting frustrated. So you decide to cheer her up by flubbing the next question. She says, 'Your turn: What's 68 + 57?' You think and then sputter, 'It's ... uh ... 126' The answer is incorrect, and you knew it all along. But given your aim this did not provide you with a reason not to give it. In the normal case you have some reason to speak the truth. But sometimes you have every reason to assert a falsehood, and when you do the practical

[14]The two objections seem to operate at different levels, and may come to the same, though I do not dwell on questions of the individuation of objections here.

[15]Here Ginsborg (2010, 5) summarises the view of Hattiangadi (2006, 227, 2007, 186–188). Also see Verheggen 2011. It might be wondered whether fiction is a similar counterexample (if lying is a counterexample at all). The answer depends on whether fiction involves genuine assertion or pretend assertion, a discussion I wish to avoid here. (With the exception of Currie (1986) and Parsons (1980), the consensus is that fiction is best understood in terms of pretending to engage in ordinary speech acts. See Walton 1990; Currie 1990; Martinich 2001; J. Searle 1979. Hoffman (2004) persuasively argues that Walton's view can be understood in terms of speech acts, even though Walton himself explicitly denies it.)

[16]Whiting 2007, 137, 2009.

valence of the claim of correctness is reversed, which is just to say that it has no valence of its own.

A proposition's mere falsehood is not even a *pro tanto* reason against asserting it. The meaning of "cat" does not entail a norm like (D) because there are no such norms—not even defeasible ones.[17]

3.2 | The assertion objection

But even if there are norms like (D), their existence does not itself show that meaning is normative. Such norms might show something not about the nature of *meaning*, but the nature of *assertion*.[18] Put another way, a norm like (D) might exist, but it might not be a *semantic* norm (i.e., one directly entailed by a meaning fact). It might instead be some kind of *assertoric* norm. If the relevant meaning fact were purely descriptive, there might still exist a specific norm for assertion, like (D), so long as there also exists the right general, constitutive norm for assertion—e.g., "only assert the truth." If so, meaning is not normative, assertion is.[19] If normativity is merely a feature of assertion, then theories of assertion need to explain it, but theories of meaning do not.[20]

4 | AUTHORITY-CONFERRING SEMANTIC NORMS

The two previous sections presented relevant background information on the normativity of meaning. Section 1 was an extremely brief presentation of the normativist thesis, and section 2 was an extremely brief presentation of the main arguments against that thesis. The present section—section 3—begins with more background information, though this time from analytical jurisprudence. By the time we reach subsection 3.4, we will have done all the stage-setting required to present the core claim of this paper, that semantic norms might be thought of as having a different structure than they have so far been thought to have.

[17]The purpose of this section is merely to lay out the objection, not to endorse it. (*Mutatis mutandis* for the following section about assertion.) It may seem to some that Rosen's example is merely a denial of Whiting's position and not an argument against it. For our purposes, however, what matters is that the new version of normativism introduced in the present paper will side-step this not-uncommon objection altogether.

[18]Wikforss 2001, 206–207; Hattiangadi 2007; Glüer and Wikforss 2009. Also, Boghossian (2003, 38) offers a response to the corresponding objection to normativism about mental content.

[19]Norms similar to (D) might exist for other kids of speech acts—questions, commands, etc. But, similarly, this may be not because meaning is normative, but because the other speech acts are.

[20]The assertion objection can either be formulated at the level of semantic norms, or at the level of semantic correctness. In the literature on the normativity of meaning, a distinction is drawn between (a) meaning facts entailing facts about which uses of linguistic expressions are correct/incorrect and (b) meaning facts entailing facts about how speakers "ought" or "ought not" use linguistic expressions. Most anti-normativists (the exception being Bilgrami 1993) accept (a) but deny (b). In construing the assertion objection as an objection to (a), it would be pointed out that merely forming a false proposition in the antecedent of a conditional is in no way "incorrect". So, it would be argued, correctness facts are fundamentally facts about speech acts and not about meaning itself.

4.1 | Three types of norms

Consider again the semantic norms discussed in section 1:

(A) "cat" ought to be applied to all and only cats.
(B) "cat" ought to be applied only to cats.
(C) *S* ought to apply "cat" only to cats.
(D) *S* ought to assert "That's a cat." only of cats.[21]

All of these norms are *duty-imposing*.[22] They specify what individuals must or must not do. They impose duties or obligations.[23] Similar norms are *permission-granting*. They specify what duties or obligations individuals *lack*. But there are norms of a third type, which we can call *authority-conferring*. I will suggest that semantic norms might be thought of as instances of this third type. But in order for that suggestion to even make sense to those without a background in the philosophy of law, we will need to spend just a few pages further explaining how these three types of norms are being understood.

Duty-imposing norms can be formulated as follows:

Duty-Imposing: *S* ought to φ.

The "ought" here need not be an all-things-considered "ought." This is also the case for permission-granting norms, which can be formulated as follows:

Permission-Granting: ~(*S* ought to ~φ).

To say that someone is permitted to do something is to say that they lack a duty or obligation not to do it.[24] Of course, being permitted *in some sense* to do something is equivalent only to lacking a duty or obligation *in that sense* not to do it. For example, if a child is permitted, by the rules of her parents' household, to stay out of the home until 10 pm, then she does not have a duty or obligation, within her parents' system of household rules, to return home earlier than 10 pm. She lacks that duty. But she may have *other* duties to return home before 10 pm, perhaps from a promise she made, a city-wide curfew, etc.

As we can see, duty-imposing and permission-granting norms are distinct, but closely related. Permission-granting norms do not impose duties. They remove duties, or they reiterate that the relevant duties were never imposed in this first place. However, they do directly concern duties. It is for this reason that H.L.A. Hart groups them together with duty-imposing norms. In Chapter

[21]And (D') from the footnote at the end of section 1.

[22]It might be argued that (A) and (B) do not really impose duties because they are not agent-implicating. Indeed, the fact that they are not agent-implicating seems to suggest that (A) and (B) are not norms at all. As we will see, this complication is irrelevant to the central point made in the remainder of this section because if they constitute norms at all, then they fall squarely on the duty side of the relevant distinction.

[23]The terms "duty" and "obligation" are used synonymously throughout the paper.

[24]In these ways, a permission is different from a right. To be permitted to do something is to lack a certain kind of duty not to do it. To have a right to do something is for *others* to have a certain kind of duty not to interfere, in some way or another, with one's doing it.

3 of *The Concept of Law*, Hart introduces something that almost immediately became part of the orthodoxy of 20th- and 21st-century philosophy of law: the distinction between duty-imposing and authority-conferring norms.[25] He introduced this distinction as part of his attack on John Austin's theory of law. Austin thought that legal systems could be understood as a series of commands backed by sanctions. This theory has some initial plausibility when applied to duty-imposing laws. For example, a law requiring citizens to pay taxes imposes a duty on those citizens. This law can somewhat plausibly be thought of as a command, issued by a sovereign, directed at citizens, and backed by the threat of sanction. But Hart points out that there are other laws that simply cannot be understood on this command-based model. A law that says that a governor can pardon convicted criminals does not say what the governor must or must not do. Rather, it authorizes her to do something: pardon criminals. If the governor decides to pardon no one or everyone, she may thereby violate *some* duty that she has—perhaps a moral duty or some other legal duty—but she does not violate a duty imposed by the law authorizing her to pardon criminals. That law is simply not in the business of imposing a duty. It is in the business of giving the governor *authority*. Hart's point is that authority-conferring laws cannot be reduced to duty-imposing ones, and that since Austin's command-based theory cannot accommodate such laws, it fails.

But what exactly is "authority" in this sense? Authority is a type of power or ability. Specifically, it is normative power. To have the authority to do something is to be able to do that thing in virtue of and within the context of some system of norms. A bodybuilder who has the physical strength to lift a heavy weight has that power in virtue of the nature of her muscles and bones. A governor who has the authority to pardon criminals has that power in virtue of the nature of the legal system within which she holds the office of governor. It is in this sense that authority is *normative* power.

How exactly this kind of normative power works is a difficult and substantive question, on which we need not commit to an answer for the purposes of this paper. But Hart himself has a view of how this kind of normative power works. And for the purpose of illustration, it is helpful to briefly lay it out. Hart says that authority-conferring norms—the norms that generate and maintain normative power of this sort—are "recipes for creating duties."[26] The idea is that a law authorizing a governor to pardon criminals gives that governor the power to do

[25]Hart 1961, 33. Hart used "power-conferring rules" to refer to what I call "authority-conferring norm." Hart emphasized that he was using "power" to mean *authority* or *normative power*, not mere *physical ability* or *non-normative power*, and I have switched to "authority" to avoid precisely this confusion. Some philosophers impose a distinction between "rules" and "norms," but nothing discussed here hinges on any of the common ways of distinguishing the two, and so I understand them as synonymous. Also, as I am using the terms, a law is a kind of norm.

[26]Hart 1961, 33. Hart does not provide a full-blown theory of authority or normative power, and he does not need to. He provides enough of a characterization of his understanding of the phenomenon to make his argument against Austin clear and, frankly, decisive. Two things are worth mentioning about authority as Hart understood it. The first is that the fact that one has some authority does not entail that others *have* certain duties. (Authority is, in this way, distinct from rights, the existence of which may plausibly entail that duties had by others. Though, in *The Concept of Law*, Hart occasionally says "right" when I believe he means to say "authority." But that is another matter.) Rather, it entails that one has the capacity to crefate duties for others. Merely having some authority does not make it the case that others have certain specific duties. It is only when one *exercises* that authority that others come to have certain specific duties. Second, though Hart only characterises authority-conferring norms as recipes for creating duties, there is nothing he says that excludes from his conception that such norms are also recipes for creating other kinds of authority. So, e.g., my authority to write a legally binding will may be a recipe for creating duties for others as to how my estate is to be distributed, but it is also a recipe for creating the authority that my executor will have to decide how to best distribute my estate in keeping with the guidelines provided in the will.

something that others cannot do, and that power can be understood in terms of her ability to generate duties that apply to the conduct of other legal officials. When the governor exercises her power by actually pardoning someone, that act imposes a duty on officials of the judicial system to—let us say, for the sake of example—release the former criminal, update the relevant state records, and so on.[27]

If Hart's conception of the nature of authority is right, then both authority-conferring and duty-imposing norms do, in some sense, generate duties. But the *ways* that they do so differ in two significant respects. First, the duties that result from authority-conferring norms are imposed *indirectly*. When a law authorizing the governor to pardon criminals is enacted, officials of the judicial system do not yet have the relevant duties—at least not the specific ones—requiring them to release the pardoned individuals, update state records, and so on. Rather, those duties are only imposed if and when the governor decides to exercise her authority by actually pardoning someone. By contrast, when a duty-imposing law requiring citizens to pay taxes is enacted, the specific duty is imposed directly, simply by virtue of the law's enactment. Second, the duties generated by authority-conferring norms are not imposed *on the agent who has the authority*.[28] That is, whereas a duty-imposing law requiring citizens to pay taxes imposes a duty *on those citizens*, an authority-conferring law giving a governor the authority to pardon criminals does not itself impose duties *on the governor*.[29]

Again, for the purposes of this paper, we do not need to commit ourselves to Hart's particular conception of authority as the power to generate duties for others. Hart's particular conception is merely included here as a way of illustrating what it might mean for some power to be "normative." The more general, important point is this: authority is the normative power or ability to do something within the context of a system of norms.

Normative power of this sort is ubiquitous. To use another example of Hart's, a law specifying how a citizen writes a legally binding last will and testament is authority-conferring. Citizens have no legal duty to write or not write a will. Rather, they have the authority—the legal power—to do so. A citizen who attempts to write a will but forgets to sign it, for example, has not violated a duty and she faces no sanction. Rather, her attempted will is simply null and void.[30] She has the authority to write a will, and she intends to write one, but she fails, just as a bodybuilder who is strong enough to lift a weight might attempt to lift it, but fail, perhaps because she does not plant her feet properly. Alternatively, someone may fail to exercise some authority because they lack that authority altogether. An ordinary citizen proclaims "I hereby pardon so-and-so," but nothing happens because she lacks the authority to pardon criminals, just as someone might fail to lift a heavy weight because she lacks the physical strength.

Are authority-conferring norms genuinely distinct from permission-granting ones? Yes. Whether someone has permission to do something is a matter of whether or not they have a duty to refrain from doing *a thing that they are otherwise capable of doing*. The child who is permitted to be out of

[27]From the fact that authority-conferring norms generate duties it does *not* follow that they are the *only* way that duties are generated, or that *all* duties or permissions are backed by authority.

[28]This is why the potential reducibility of authority to duties does not obviate the distinction. This is discussed a bit more in section 3.4.

[29]See Hart 1961, 26–49.

[30]One might wonder if the relevant law really does impose a duty on citizens to write a will and the sanction for failing to write a will is simply nullity. This is the main objection that Hart considers in Chapter 3 of *The Concept of Law*, and he is widely taken to have been successful in his response.

the house before 10 pm is, we can assume, entirely capable of opening the door, exiting the house, and walking or driving to some destination other than the house itself. But whether someone has the authority to do something is a matter of whether or not there exists a normative system empowering them to do *a thing that cannot be done outside that normative system*. Pardoning criminals is not something that a governor is otherwise capable of doing. It is not as if any ordinary citizen can pardon criminals, but that they are prohibited from doing so, leaving the governor as the only one with the permission to pardon. Rather, ordinary citizens are *unable* to pardon. The law specifying that the governor can pardon criminals does not give the governor permission to do something that she is otherwise able to do. It empowers her—giving her an ability that can only be had within the relevant system of norms—to do it.

4.2 | The structure of authority-conferring norms

Authority-conferring norms can be put in the following form:

> **Authority-Conferring:** *S* is authorized to φ; to exercise this authority *S* ψs.

In an authority-conferring norm, φ is an action *described in the characteristic vocabulary of the relevant social institution*, and ψ is an action described in either institutional or neutral vocabulary. For instance, a law for creating a will might be formulated as: "Citizens and legal residents of sound mind and over the age of 18 are authorized to create a last will and testament; to exercise this authority the citizen or legal resident signs such-and-such a document in the presence of two witnesses." Here φ is "create a last will and testament," which is a legal status characterized in legal language. The action ψ is "sign such-and-such a document in the presence of two witnesses," which is a set of instructions for achieving that legal status.[31]

One thing to note here is that, like permission-granting norms, authority-conferring norms are not the kind of thing that can be *violated*. A duty-imposing norm specifying that citizens must pay taxes is violated when a citizen fails to pay taxes. But a permission-granting norm specifying that a child does not have a duty to stay at home during a certain period of time is violated neither when the child stays home nor when she goes out. It does not make sense to say that one's permission has been "violated," because permission is not duty; it is the absence of duty. Similarly, an authority-conferring norm specifying that a governor can pardon criminals is violated neither when she pardons someone nor when she pardons no one. An authority-conferring norm specifying that certain citizens can write legally binding wills is violated neither when a citizen writes a will nor when a citizen does not write a will.

At first glance, one may think that authority-conferring norms *can* be violated, and that that violation occurs not when *S* fails to φ, but rather when *S* attempts and fails to ψ. But as mentioned in the previous subsection, this is not so. If someone attempts to write a will but does not sign it in front of two witnesses, then she has not violated estate law. Rather, when she eventually dies, she simply dies intestate—without a valid will. And, as Hart points out, nullity is not a sanction.[32]

[31] In actual legal systems, the details of ψ are sometimes made explicit and other times left to be fixed by precedent.

[32] Hart 1961, 33–35.

Of course, failing to write a legally valid will can violate *some* duty-imposing norm—such as a norm of one's condominium board requiring all homeowners to have a valid will. Or if one intends to write a will, then, perhaps, failing to do so violates one's intention—violates, so to speak, a duty to oneself. But the crucial point is that failing to write a legally valid will does not violate the law.[33] And that is because the law, in this case, is authority-conferring and is not the sort of thing that can be violated.[34]

It is worth clarifying that both duty-imposing and authority-conferring norms can be *constitutive* of practices, but only duty-imposing norms can be *regulative* of previously possible behavior. A duty-imposing law requiring pedestrians to cross only at crosswalks is a constitutive norm in the sense that there are new forms of behavior made possible by the existence of the law—e.g., jaywalking. And this same law is also a regulative norm in the sense that it regulates a previously possible form of behavior: crossing the street. By contrast, an authority-conferring municipal bylaw empowering the town council to make new traffic laws is also a constitutive norm in the sense that there are new forms of behavior made possible by its existence—e.g., enacting a new traffic law. But it is not a regulative norm because there is no description of any kind of behavior that was possible prior to the existence of the law and that that authority-conferring municipal bylaw can naturally be said to "regulate."[35]

4.3 | Authority outside the law

So far, only two examples of authority-conferring norms have been mentioned, and both were *legal* examples. Do authority-conferring norms only appear in the law? No. They are found in all moderately sophisticated, human-created normative systems. The obvious examples related to games involve umpires or referees. A baseball umpire may have a duty to call balls and strikes. But she also has the authority to do so. If a player or fan screams "Strike one!," nothing happens. There is no change to the batter's in-game normative status. But an umpire's pronouncements do change the normative state of affairs, and that is because there is some rule granting an umpire the relevant authority.

But authority-conferring norms are even more fundamental to games than this. For example, what is the rule of chess regarding the movement of a rook? The rule is not that a rook must be moved vertically or horizontally, because it is not true that a rook must be moved at

[33]This does not entail that some kinds of authority do not come attached to duties. An instructor has the authority to issue grades for a course. But she also has a duty regarding how she is to exercise that authority—e.g., fairly and impartially, among other things. When she grades unfairly her duty has been violated. Her authority has not been violated; it has been *abused*.

[34]This point about authority-conferring norms not being subject to violation, supplies some of the best evidence that philosophers of mind and language simply have not considered that semantic norms might be authority-conferring. We can see this from the fact that several anti-normativists rely on the thought that all norms can be violated. See Bilgrami 1993; Wikforss 2001, 213. Hattiangadi (2006, 221) says that normative statements "tell us what to do, whereas non-normative statements simply describe how things are." If authority-conferring norms are normative, then it is not true of all normative statements that they 'tell us what to do'. Some of them give us authority to tell others what to do.

[35]See Rawls 1955; J. Searle 1969; Raz 1975; Ludwig 2017.

all.[36] In *Practical Reason and Norms*, Joseph Raz points out that the rule of chess regarding the movement of a rook is authority-conferring.[37] A player has the authority to move her own rooks vertically or horizontally. And she lacks the authority to move a rook diagonally, just as she lacks the authority to move any of her opponent's pieces. This is how it is possible to make an incorrect move when attempting to play chess. One makes a move without the relevant authority.[38]

4.4 | Semantic authority

The central idea of this paper is that semantic norms can plausibly be understood not as duty-imposing, but as authority-conferring.[39] There are many different ways to work out this proposal. The purpose of this paper is not to argue for any one of them in particular. But, nonetheless, in order to show how understanding semantic norms as authority-conferring allows normativism to escape the main objections to it, it is helpful to work with a specific putative semantic norm. Consider, therefore, the version of normativism according to which the fact that "cat" means *cat* in the language of a speaker *S* entails that:

(E) *S* is authorized to refer to cats; to exercise this authority *S* says "cat."

The most important thing to explain about (E) is what it is to be *authorized to refer to cats*. But before launching into that explanation, it is worth making five smaller points. First, (E) is a norm, not a definition. As such, it should not be expected that saying "cat" is necessary and sufficient for

[36]Nor is the rule that a rook must not be moved diagonally, as one who moves a rook diagonally is not playing chess (because she does not want to play, does not know how to play, etc.) and is not violating a rule of chess. Of course, moving a rook diagonally can be a violation of a duty, such as the kind one acquires if one promises to play chess. But this duty is promissory, and not constitutive of chess.

[37]Raz 1975, 115.

[38]What about cheating? One might try to understand moving a rook diagonally with the intention of having one's opponent not notice as a violation of a duty-imposing rule against moving one's rook diagonally. But it is better, I suggest, to understand one as having the authority to move a rook only vertically or horizontally as well as having a general duty only to exercise the authority that one has within the game. That is, when one cheats one is violating something, but what is violated is one's opponent's trust—or a norm requiring one to only exercise one's legitimate authority or to, so to speak, play by the rules—and not a specific rule regarding movement of a rook.

[39]We might wonder about the alternative possibility that semantic norms are permission-granting. I consider that question to be outside the scope of this paper, though it is discussed elsewhere. See Hattiangadi 2009, 58; Whiting 2009, 544–545. Also, we might wonder whether authority-conferring norms are really norms and whether the view that they are directly entailed by meaning should count as normativism about meaning. But the labels "norm" and "normative" are not what ultimately matter. Whether or not we are willing to call authority-conferring norms "norms," what matters is whether they present a constraint on our theories of meaning in the way that the normativity of meaning has been thought to. And they do. In order for an authority-conferring norm to be in place, there must be a normative system—one consisting of duties—in place. So authority-conferring norms are just as difficult for theories of meaning to explain and duty-imposing norms are.

referring to cats. And, indeed, it is neither.[40] Second, the corresponding norms for other subsentential expressions (e.g., quantifiers, adjectives, connectives, modal auxiliaries) would be very different from (E). They may, for example, not involve reference. Indeed, the relevant norm for referring expressions such as "cat" may be different from (E). The purpose of mentioning (E) is merely to give us an example with which to illustrate some of the results of thinking of semantic norms as authority-conferring. Third, norms like (E) confer the authority *to use a word meaningfully*, not to give a word some standing meaning. The fact that "cat" means *cat* in a language entails that speakers of that language have the authority to use "cat" to refer to cats, not the authority to establish *cat* as the meaning of the word "cat" (though depending on *S*'s social position, she may have that authority as well). Fourth, there is nothing illicit about the fact that the statement of the semantic norm (E) includes "refer," which is a semantic term.[41] As we saw in section 3.2, the first action mentioned in all authority-conferring norms is characterized in the status-laden terminology of the relevant social institution.[42] Fifth, the kind of reference that would be involved in a semantic norm like this would have to be linguistic interpersonally communicative reference. Private mental reference may or may not be closely related.[43] But the focus of this paper is linguistic meaning—meaning that occurs within the social institution of language.

Returning to the larger question about (E), what exactly is it to be "authorized to refer to cats?" Who has this authority and who lacks it? And what happens if someone refers to cats, or attempts to refer to cats, when she lacks the authority to do so? Like a governor's authority to pardon criminals or a chess player's authority to move their rook, the authority to refer to cats is a *normative power*—it is an ability that one possesses because of the existence of a system of rules and one's place within that system. To get a better sense of what this amounts to, consider two different ways that someone can lack the authority to move a chess piece, such as a rook. First, one may not be the right player, or any player at all. A chess player lacks the authority to move her opponent's pieces.[44] If she reaches over and moves her opponent's rook to another square on the board, her action does not count as a move in the game. Of course, she is *physically* able to move the circular wooden icon that resembles a castle or turret. But because she lacks the relevant authority, when she moves that wooden icon, the normative status of the pieces and players in the game do not change. When she moves that wooden icon it does not *count* as moving the rook.

[40]A speaker can say "cat" without referring to cats, for instance, when practicing pronunciation. And there are many other ways to refer to cats, for instance, with indexicals or by saying "gato". The same is true of many, perhaps most, authority-conferring norms, including the norm granting a governor the authority to pardon criminals, the norms granting chess players' authority for the movement of various pieces, etc. Also, we may wonder how norms involving the reference of individual words relate to Frege's context principle. This is discussed in a note at the end of section 3.6.

[41]Similarly, norms for other parts of speech will include other semantic terms such as the verb "to predicate".

[42]By analogy, in computer chess, a player is authorized to move her rook to a6; to exercise this authority she clicks on the square marked "a2" and then on the square marked "a6." The phrase "move her rook to a6" is phrased as a status in the game of chess, whereas "click on the square marked 'a2' and then on the square marked 'a6'" is phrased neutrally.

[43]One of the closest theorized relations, familiar from J. A. Fodor 1975; J. R. Searle 1980; Sayre 1986; J. A. Fodor 2008; J. Fodor 2009, is that linguistic reference is derivative from private mental reference or intentionality. For one alternative view, see Dennett 1989; Newton 1992.

[44]As Raz 1975 convincingly argues, a player is not prohibited or forbidden from moving her opponent's pieces, but rather lacks the authority to do so.

Similarly, if two people are playing chess in a park and someone jogging past reaches down and moves one of the pieces on the board, then that move does not count because the jogger lacks the relevant authority. Once she jogs away, the players can simply return the out-of-place icon to its previous location and carry on with the game.

A second way for someone to lack the authority to move a chess piece is for there to be no game going on. If someone walks past an unattended chess board and moves a rook, then her action does not count as making a chess move. She lacks the relevant authority—the normative power to move a rook—because outside the context of a game of chess *no one* has that authority.[45] It may or may not be possible to play chess alone, against oneself. If it is possible, then, in that circumstance, the player has the authority to move any of the rooks on the board (though only at certain times). That is because in such a circumstance there is an ongoing game of chess. But when there is no game of chess, there is no authority to move a rook.

Turning to the linguistic case, can one lack the authority to refer to cats in the same two ways? It seems to me that one can lack the authority to refer to cats in one of these ways, but not the other. First, just as one lacks the authority to move a rook when there is no game of chess, so too one lacks the authority to refer to cats when the relevant system of norms within which that authority would be embedded—the language—does not exist. A game of chess lasts only a few minutes or hours. But a system of linguistic norms is perhaps best thought of as existing for decades or centuries, with the rules slowly changing "mid-game," so to speak. Suppose that neither the English language nor any language remotely resembling it exists. A person opens her mouth and utters a wholly novel string of phonemes that sounds like "cat." In this case, it seems to me that this person has failed to *thereby* refer to cats. She lacks the authority to do so because the relevant system of rules within which that authority would have to be embedded does not exist. She may be able to refer to cats in *other* ways—for instance, in another language. Or, depending on how one thinks about one of the more contentious issues in the philosophy of mind and language, she may or may not be able to have private intentional mental states directed at cats. And such states, if she happens to have any, may or may not be closely related to meaningful public linguistic utterances.[46] But our focus here is on the latter. And the thought is that when the relevant language does not exist, the relevant authority—the ability to linguistically refer to cats—does not exist either.

Also contentiously, just as with a game of chess, it may or may not be possible for someone to have a personal language. (I do not mean a Wittgensteinian private language in which the objects of linguistic reference are *in principle* only knowable to a single subject. Rather, I mean a personal language in which a single speaker and no one else refers to ordinary, publicly observable objects of linguistic reference.) If such a language is possible, then someone speaking that language may have the authority to refer to cats. But that is because the relevant system of rules exists and her authority is part of that system. Just as one may lack the authority to move a rook when no game of chess is ongoing, so too, it seems to me, one may lack the authority to refer to cats when the system of rules in which that authority would be embedded does not exist. In such

[45]Similarly, someone in a Hobbesian state of nature lacks the authority to pardon criminals. It is not merely that she is not a governor authorized to pardon criminals, though she is not. Rather, in this circumstance, no one can pardon criminals because (a) criminality is a legal status that does not exist in a circumstance where there are no laws and (b) no one has any legal authority in such a circumstance.

[46]I have in mind here views, mentioned in a previous note, such as those of J. A. Fodor 1975; J. R. Searle 1980; Sayre 1986; J. A. Fodor 2008; J. Fodor 2009, according to which linguistic reference or intentionality is derivative of private mental reference or intentionality.

a case, it is not quite correct to say that one refers to cats even though one lacks the authority to do so. Rather, because one lacks the authority, one is *unable* to refer to cats (at least with the English word "cat.").[47]

But there was a *second* way for someone to lack the authority to move a rook—not being the right player. Is there a linguistic analogue to this? I do not believe so.[48] Unlike the legal authority to pardon criminals or the chess-based authority to move a rook, language, interestingly, seems to be a fundamentally egalitarian normative system. Of course, there may be conditions that a putative speaker must meet in order to successfully use "cat" to refer to cats. But, roughly speaking, anyone can refer. Ordinary citizens lack the normative power to pardon criminals. Chess players lack the normative power to move their opponent's pieces. But there is no class of individuals—at least not one that I can think of—who lack the normative power to refer to cats. Of course, a speaker may be *forbidden* from referring to cats. But the fact that someone is legally, societally, or in some other sense prohibited from referring to cats does not entail that she lacks the linguistic authority—the normative semantic power—to do so. Indeed, it only makes sense to prohibit this speaker from referring to cats if she is *able* to refer to cats. It is only necessary to impose such a prohibition if, were she to say "cat," she would succeed in referring to cats.

Does this fact about linguistic authority—that it is, at least in some instances, widely possessed—suggest that there is something suspect about (E) or about the general claim that semantic norms are authority-conferring? I do not think so. Though it is typically scarce, there is nothing about authority *per se* that entails that it must be restricted to a select few. We can easily imagine cases—and, indeed, some such cases may be actual—where types of legal authority are had by all. Suppose, for instance, that there is a legal system in which anyone—lawyer or layperson, citizen or foreigner, human being, or sentient alien life form—has the standing to file a claim in court. Does the maximally inclusive nature of this legal authority suggest that it is not properly understood as *authority*? No. It is not a mere permission.[49] It is a normative power. It is an ability that people have in virtue of the existence of the legal system and their place within it. The point is just that, like in the linguistic case, there are no restrictions on who has that place within the system. This fits well with the structure given to authority-conferring norms in section 3.2. Norms like (E) have two parts: one specifying who has the authority and another specifying how that authority is exercised. In the case of

[47]This entire way of thinking about linguistic reference has a Wittegenstian flavor to it, and it is thoroughly contentious, both in aggregate and in the details. The point is not to defend or even present, at least not in any detailed way, a broad, sweeping view of language of this sort. Rather, the point is just to give some illustration of what is meant by "the authority to refer."

[48]I do not think the central claim of this paper depends on giving this answer. Indeed, if there were a case in which one person had the authority to refer to cats, but someone else lacked it, then that might be good for the claim that semantic norms are authority-conferring. The point of this discussion is simply to illustrate what the authority to refer is meant to be.

[49]We should also not be misled into confusing the authority to file a claim with the right to file a claim. In most cases, the two will come together. But they are importantly different. A right corresponds to a set of duties had by others requiring them to enable or not interfere, in some way or another, with the behavior of the person who holds the right. If I have a right to file a claim, then various agents of the state—e.g., judges, bailiffs, and lawyers—have duties to process my claim. And others—perhaps all my fellow citizens—may have a duty not to impede my attempts to file a claim. But the fact that I am authorized to file a claim is just the fact that my submitting a certain piece of paper *counts* as filing a claim.

semantic norms, plausibly, the first part is maximally broad or inclusive and the second part is more restrictive.[50]

We are now in a position to answer the questions posed earlier. What is it to be "authorized to refer to cats"? It is to have a status within a system of linguistic norms such that one's utterances of "cat" *count* as references to cats. Who has this authority and who lacks it? It appears that everyone has it, and the main way to lack it is for the system of linguistic norms not to exist. What happens if someone refers to cats, or attempts to refer to cats, when she lacks the authority to do so? It is not possible to refer—at least not to linguistically refer—to cats without the authority to do so, just as it is not possible to pardon a convicted criminal without the authority to do so. If someone lacks the authority, but attempts nonetheless, then their utterance fails to count as a reference to cats.

This is only a rough sketch of what it would mean for semantic norms to be authority-conferring. But the goal of this paper is merely to introduce a promising direction for the normativity of meaning debate. The question is: what normative facts, if any, are entailed by meaning facts? It is a question of how meaning relates to language use. Does the meaningfulness of a speaker's expressions entail that she has an obligation to do something or the power to do something? The suggestion is that it entails that she has a power. Sometimes, when a person opens her mouth, she merely emits sounds. But other times, when she opens her mouth her sounds are meaningful expressions within a language. The adherent of duty-based normativism claims that the fact that an expression is meaningful directly entails something about when and if a speaker ought to utter it. The adherent of authority-based normativism claims that the fact that an expression is meaningful directly entails something about what a speaker is able to do by uttering it. My suggestion is that this latter view has at least some initial plausibility.

4.5 | Sentential and sub-sentential norms

One worry that is worth addressing immediately is that (E) does not seem analogous to (D). The difference between them was supposed to be that (D) is duty-imposing while (E) is authority-conferring. But they also seem to differ in a second respect: (D) is a sentential norm governing the use of "That's a cat," whereas (E) is a subsentential norm governing the use of just the single word "cat." And as we will see below, this sentential/subsentential difference is what allows (E) to escape at least one of the objections to (D) discussed in section 2. The worry is that any advantages that (E) has over (D) will be due not to its authority-conferring structure, but to the fact that it is subsentential.

This need not worry us, however, because (E) can be formulated as subsentential *because it is authority-conferring*. So any objections that it avoids by being subsentential are ultimately the result of its authority-conferring structure.

To see this, recall that in section 1 the adoption of the sentential norm (D) resulted from the inadequacy of the subsentential norm (C), which was "*S* ought to apply 'cat' only to *cats*." The

[50]Related to this, we might wonder about the possibility of the incorrect use of language on this picture. There are many ways to use a linguistic expression incorrectly, and space limitations prevent me from discussing them all. But one way is that a speaker might mistakenly assert a falsehood, such as by asserting that something is a cat when it is really a dog. In this case, *vis-à-vis* the semantic authority-conferring norm (E), the speaker has actually *succeeded*. That is, she referred to cats (or she referred to the universal of cathood, or she predicated of something that it is a cat, etc.). Of course, such a speaker may still have violated some other norm, such as a truth-telling norm. But whether that is the case depends on whether we have the intuitions discussed in section 2.1.

problem with (C) was that any norm prohibiting the *mere application* of "cat" to non-cats was far too strong. So to make (C) plausible, it was reinterpreted or reformulated so as to govern only assertion. But while *predicates* can be applied, only *full sentences* can be asserted. So to reformulate (C) to govern only assertion, it had to be made sentential. That is how we got from the subsentential norms (A), (B), and (C) to the sentential norm (D). But when considering authority-conferring semantic norms we are able to remain at the subsentential level. The authority-conferring norm (E) is analogous to (C), and when thinking of semantic norms as structured like (E) we are not forced to ascend to the sentential level because authority-conferring semantic norms do not mention the "application" of a term.

All of this may seem rather abstract at the moment, but its relevance becomes clear in section 4. What matters is simply that any advantages that accrue to (E) because it is subsentential (e.g., that it avoids the assertion objection because it involves *reference* as opposed to *assertion*) are ultimately due to its being authority-conferring.[51]

5 | RETURNING TO THE OBJECTIONS TO NORMATIVISM

Understanding semantic norms as authority-conferring allows normativism to escape the objections from section 2.

5.1 | Avoiding the permissibility of lying objection

The permissibility of lying objection rests on the following intuition: sometimes lying does not violate any norm. Therefore, there cannot be norms like (D) that apply to all meaningful language use. However, as mentioned earlier, authority-conferring norms are not the kind of thing that can be *violated*. So while the fact that some lies do not violate any norm is evidence that there are no duty-imposing semantic norms, it is not evidence that there are no authority-conferring semantic norms. Since authority-conferring norms are not the kind of thing that can be violated, their existence is compatible with any intuitions whatsoever about when norms are or are not violated.[52]

[51]One might also worry that any talk of subsentential norms in the first place violates Frege's context principle, since it seems to involve making sense of the word "cat" referring to *cats* outside of the context of that word's use within a whole sentence. Luckily, however, talk of subsentential linguistic norms—including the talk of them that we engaged in when discussing (A), (B), and (C)—does not require that we reject the context principle. It is possible to have the authority to do something that can only be done within the context of some other rule-governed activity, which is how we can understand reference if we wish to simultaneously entertain the authority-based understanding of normativism and Frege's context principle.

[52]This is an exaggeration. There is at least one obvious way that the existence of an authority-conferring norm might be incompatible with some intuitions about when norms are violated: if the authority-conferring norm is reducible to duties had by others, then the existence of the norm might be threatened by intuitions about when those duties had by others are violated. Relevantly for our purposes, though, those cannot be intuitions about whether the *speaker* violates a norm by lying, as the permissibility of lying objection holds, because, as noted earlier, authority-conferring semantic norms that apply to a speaker cannot be reduced to duties had by *her*.

The fact that so many in the normativity of meaning debate—e.g., Wikforss (2001, 213) and Bilgrami (1993)—assume that the possibility of violation is required for normativity is evidence that they have simply failed to consider the possibility of authority-conferring norms.

5.2 | Avoiding the assertion objection

Authority-conferring norms like (E) are subsentential, and make no mention of particular speech acts, like assertion. So, as most readers have already gathered, there is no danger that their normativity is a feature of the nature of a particular speech act as opposed to the nature of meaning itself. We know that (E) is a semantic (and not assertoric) norm because it makes no mention of assertion. And, unlike (C), since (E) does not concern the mere *application* of terms, there is no need to ascend to a new norm at the level of sentences or speech acts.[53]

6 | CONCLUSION

This paper proposes a revision of the normativist thesis about meaning. And though it does not constitute a wholesale defense of normativism, the proposed revision does allow normativism to avoid the most significant and frequently-cited objections to it. Thinking of semantic norms, within the normativity of meaning debate, as authority-conferring represents a departure from the entire history of that debate, going all the way back to Kripke's original statement of the normativity of meaning:

> The point is not that, if I meant addition by '+', I will answer '125', but that, if I intend to accord with my meaning of '+', I *should* answer '125'.[54]

Kripke understands meaning as entailing a duty-imposing norm, though he does not say what that norm is. Since the version of normativism introduced in this paper departs from all previous forms, it should not be surprising that it saves that doctrine from its most pressing objections.

But the claim that semantic norms are authority-conferring really is only about the structure of meaning normativity, not about the nature of the normativity itself. And that is its greatest

[53]Related to the issue of sentential and subsentential expressions, it might be worried that reference is only possible within the context of sentential expressions, or perhaps even speech acts. This, of course, is controversial and I do not wish to take a stance on it here, but even supposing it is true this does not count against the suggestion that semantic norms are authority-conferring. Consider, by analogy, an authority-conferring law specifying how citizens can designate the executor of a will. All citizens above 18 years of age can designate an executor, but doing so only has legal force within the context of a will. If there is no associated will, then specifying who is responsible for executing the will is vacuous. Nonetheless, the authority to specify an executor is one that citizens have and it should be understood as a form of authority, not a duty or permission (though, citizens will occasionally also have a duty to do so and almost always have permission to do so). Analogously, referring to cats may only have effect within the context of some larger linguistic action, but this does not undermine the thought that the right way to understand that linguistic relation (i.e., reference) is in terms of authority, not duty or permission.

[54]Kripke 1982, 37. Also see other early work from both normativists (McDowell 1985, 1984, 1994; Boghossian 1989) and anti-normativists (J. A. Fodor 1990; Bilgrami 1993; Papineau 1999).

virtue. If semantic norms are authority-conferring, then the brand of normativity involved remains the same. So if some naturalistic resources are inadequate to explain a duty or obligation, then they will be equally inadequate to explain any corresponding authority or normative power. The normativity of meaning understood in terms of authority-conferring norms will be just as much of a non-trivial condition on the adequacy of theories of meaning as if it were understood in terms of duty-imposing norms. But the two most popular and pressing objections to the normativity of meaning thesis are both *structural* objections, concerning structural features of meaning norms: i.e., whether those norms involve implicit or explicit appeal to assertion, or whether they are the kind of norms that can be violated. It is by noticing the distinction between these two features of normativity—its structure and its, so to speak, metaphysical variety—that we sidestep the objections to normativism without diminishing its significance to our understanding of the nature of meaning.

ORCID

Jeffrey Kaplan https://orcid.org/0000-0003-1011-0126

REFERENCES

Bilgrami, A. (1993). Norms and meaning. In R. Stoecker (Ed.), *Reflecting Davidson*. De Gruyter.

Blackburn, S. (1984). The individual strikes back. *Synthese*, *58*(3), 281–301. https://doi.org/10.1007/BF00485244

Boghossian, P. A. (1989). The rule-following considerations. *Mind*, *98*(392), 507–549.

Boghossian, P. A. (2003). The normativity of content. *Philosophical Issues*, *13*(1), 31–45. https://doi.org/10.1111/1533-6077.00003

Brandom, R. B. (1994). *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard University Press.

Bykvist, K., & Hattiangadi, A. (2007). Does thought imply ought? *Analysis*, *67*(4), 277–285. https://doi.org/10.1093/analys/67.4.277

Currie, G. (1986). Works of fiction and illocutionary acts. *Philosophy and Literature*, *10*(2), 304–308. https://doi.org/10.1353/phl.1986.0046

Currie, G. (1990). *The nature of fiction*. Cambridge University Press.

Dennett, D. C. (1989). *The intentional stance*. Bradford Books.

Fodor, J. A. (1975). *The language of thought*. Harvard University Press.

Fodor, J. A. (1990). *A theory of content and other essays*. MIT Press.

Fodor, J. A. (2008). *LOT 2: The language of thought revisited*. Oxford University Press.

Fodor, J. A. (2009). Where is my mind? *London Review of Books*, *31*(3).

Gibbard, A. (2012). *Meaning and normativity*. Oxford University Press.

Ginsborg, H. (2010). Review of oughts and thoughts: Rule-following and the normativity of content, by Anandi Hattiangadi. *Mind*, *119*(476), 1175–1186. https://doi.org/10.1093/mind/fzq100

Ginsborg, H. (2012). Meaning, understanding and normativity. *Aristotelian Society Supplementary*, *86*(1), 127–146.

Ginsborg, H. (2018). Normativity and concepts. In D. Star (Ed.), *The Oxford handbook of reasons and normativity* (pp. 989–1014). Oxford University Press.

Glüer, K., & Pagin, P. (1998). Rules of meaning and practical reasoning. *Synthese*, *117*(2), 207–227.

Glüer, K., & Wikforss, Å. (2009). Against content normativity. *Mind*, *118*(469), 31–70. https://doi.org/10.1093/mind/fzn154

Glüer, K., Wikforss, Å., & Ganapini, M. (2022). The normativity of meaning and content. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Spring 2022 edition. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2022/entries/meaning-normativity/

Hart, H. L. A. (1961). *The concept of law*. Oxford University Press.

Hattiangadi, A. (2006). Is meaning normative? *Mind and Language*, *21*(2), 220–240. https://doi.org/10.1111/j.0268-1064.2006.00312.x

Hattiangadi, A. (2007). *Oughts and thoughts: Rule-following and the normativity of content*. Oxford University Press.

Hattiangadi, A. (2009). Some more thoughts on semantic oughts: A reply to daniel whiting. *Analysis*, *69*(1), 54–63. https://doi.org/10.1093/analys/ann009

Hoffman, S. (2004). Fiction as action. *Philosophia*, *31*(3), 513–529. https://doi.org/10.1007/BF02385199

Hohwy, J. (2001). Semantic primitivism and normativity. *Ratio*, *14*(1), 1–17. https://doi.org/10.1111/1467-9329.00141

Humberstone, I. L. (1971). Two sorts of 'ought's. *Analysis*, *32*(1), 8–11. https://doi.org/10.1093/analys/32.1.8

Kripke, S. A. (1982). *Wittgenstein on rules and private language*. Harvard University Press.

Ludwig, K. (2017). *From plural to institutional agency: Collective action II*. Oxford University Press.

Martinich, A. (2001). A theory of fiction. *Philosophy and Literature*, *25*(1), 96–112. https://doi.org/10.1353/phl.2001.0014

McDowell, J. (1984). Wittgenstein on following a rule. *Synthese*, *58*(3), 325–363. https://doi.org/10.1007/BF00485246

McDowell, J. (1985). Values and secondary qualities. In T. Honderich (Ed.), *Morality and objectivity* (pp. 110–129). Routledge.

McDowell, J. (1994). *Mind and world*. Harvard University Press.

Newton, N. (1992). Dennett on intrinsic intentionality. *Analysis*, *52*(1), 18–23. https://doi.org/10.1093/analys/52.1.18

Papineau, D. (1999). Normativity and judgment. *Proceedings of the Aristotelian Society*, *73*(73), 16–43.

Parsons, T. (1980). *Nonexistent objects*. Yale University Press.

Rawls, J. (1955). Two concepts of rules. *The Philosophical Review*, *64*(1), 3–32. https://doi.org/10.2307/2182230

Raz, J. (1975). *Practical reason and norms*. Hutchinson.

Rosen, G. (2001). Brandom on modality, normativity, and intentionality. *Philosophy and Phenomenological Research*, *63*(3), 611–623. https://doi.org/10.1111/j.1933-1592.2001.tb00128.x

Sayre, K. M. (1986). Intentionality and communication theory. *The Behavioral and Brain Sciences*, *9*(1), 155–165. https://doi.org/10.1017/S0140525X00021956

Searle, J. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge University Press.

Searle, J. (1979). The logical status of fictional discourse. In *Expression and meaning: Studies in the theory of speech acts* (pp. 58–75). Cambridge University Press.

Searle, J. (1980). Minds, brains, and programs. *The Behavioral and Brain Sciences*, *3*(3), 417–424. https://doi.org/10.1017/S0140525X00005756

Speaks, J. (2009). The normativity of content and 'the frege point'. *European Journal for Philosophy of Science*, *17*(3), 405–415. https://doi.org/10.1111/j.1468-0378.2009.00361.x

Verheggen, C. (2011). Semantic normativity and naturalism. *Logique et Analyse. Nouvelle Serie*, *54*(216), 553–567.

Walton, K. L. (1990). *Mimesis as make-believe: On the foundations of the representational arts*. Harvard University Press.

Whiting, D. (2007). The normativity of meaning defended. *Analysis*, *67*(294), 133–140. https://doi.org/10.1111/j.1467-8284.2007.00663.x

Whiting, D. (2009). Is meaning fraught with ought? *Pacific Philosophical Quarterly*, *90*(4), 535–555. https://doi.org/10.1111/j.1468-0114.2009.01354.x

Whiting, D. (2016). What is the normativity of meaning? *Inquiry*, *59*(3), 219–238. https://doi.org/10.1080/0020174X.2013.852132

Wikforss, Å. (2001). Semantic normativity. *Philosophical Studies*, *102*(2), 203–226.

How to cite this article: Kaplan, J. (2023). The structure of semantic norms. *Analytic Philosophy*, 64, 373–391. https://doi.org/10.1111/phib.12262