

## Gibbard on Meaning and Normativity

(to appear in book symposium in *Inquiry*)

Timothy Williamson

It is a long time since equating a sentence's meaning with its verification condition looked like a promising idea. But have we been too hasty? Let's not use the ominous phrase 'verification condition'. We can take the more fashionable normative turn, and speak instead of the condition under which a speaker *ought* to accept the sentence. If we want to be even more sophisticated, and a little bit formal, we can grade acceptance, and speak of the condition under which one ought to give the sentence a given degree of credence. We can gloss 'condition' here as 'evidential condition'. We can also emphasize how one's credence in a complex sentence ought to be related to one's credence in other sentences, especially those with which it shares constituents. As verificationism is revamped along such lines, it sounds increasingly like the position at which Allan Gibbard arrives in *Meaning and Normativity*. Here is a sample passage:

The meanings of the words in a sentence combine to explain which inferential oughts apply to the sentence and the evidential conditions under which one ought to accept or reject the sentence. A word's meaning what it does consists in the pattern of oughts that enters into such explanations. (p. 114)<sup>1</sup>

Of course, Gibbard's intention is not to perform the Vienna Circle in modern dress. His thoughtful, ingenious, impressively crafted, and subtly reflexive monograph goes far to work out an intriguing line of thought in response to Saul Kripke's suggestive remarks about the normativity of meaning in his book on Wittgenstein's rule-following considerations. And Gibbard has no difficulty in looking more carefully at normative language than the logical empiricists did. Nevertheless, the similarities in their views are genuine and extensive. Indeed, they go further than has already been indicated. For example, he highlights his reliance on an analytic–synthetic distinction:

I posit a sharp distinction between what is normative and what is naturalistic. This presumes more broadly a sharp distinction between what's synthetic and what's analytic. (p. 23)

He also assumes that our access to the external world is mediated by a phenomenal given that is epistemically transparent to us:

Phenomenal ascriptions are special in that they are always credible if true. If something is phenomenally blue then one ought to believe that it is. (p. 130)

Dogs are theoretically posited causes of certain recurrent features of our experience. (p. 133)

One response to such an account would be simply to invoke one's favourite refutation of logical empiricism, perhaps Quine's 'Two Dogmas of Empiricism'. But that would be unsatisfying. Such generic refutations rely on generic premises that have not all stood the test of time much better than logical empiricism itself did. Those with some sympathy for the project may hope that the extra layers of normative and probabilistic sophistication that Gibbard has added will suffice to protect it from the usual charges. In this discussion, I address Gibbard's account directly and in detail, although far from comprehensively. I leave it to the reader to judge how closely the problems that emerge are related to the familiar problems for logical empiricism. In the final section, I draw the moral that reliance on an obsolescent ideology of 'conceptual connections' has misled attempts to understand the links between normative thought and action, and created unnecessary obstacles to acceptance of a robustly realist and cognitivist conception of normativity, and indeed of morality.

## I

I will start with the relation of synonymy. Unlike Quine in 'Two Dogmas', I do not take it as notably obscure or problematic. It is a familiar phenomenon, recognized in linguistics, even though the underlying nature of what synonyms have in common — meaning — is a matter for theoretical investigation. One aim of Gibbard's theory is to explain what synonymy consists in.

Gibbard uses an example of synonymy across languages: 'Pierre's word "chien" is synonymous with my word "dog", and his sentence "Les chiens aboient" is roughly synonymous with my sentence "Dogs bark"' (p. 119). The reason for the qualification 'roughly' here is that 'Les chiens aboient' may differ slightly in semantic structure from 'Dogs bark', through the presence of the definite article in the former but not the latter. In such cases, if the difference in structure makes no

overall difference to the logical powers of the sentences, Gibbard classifies the sentences as analytically equivalent but not synonymous; analytic equivalence is necessary but not sufficient for synonymy. Now suppose, with Gibbard, that Pierre's sentence 'Voici un chien' is analytically equivalent to Gibbard's sentence 'Here's a dog'. Then, on his view, for the same evidential circumstance, the credence that Pierre ought to give his sentence 'Voici un chien' is the same as the credence that Gibbard ought to give his sentence 'Here's a dog'.<sup>2</sup>

One might think that Pierre and Gibbard cannot be in the same evidential circumstance, because Pierre has been exposed to evidence about French while Gibbard has been exposed to evidence about English. To finesse that issue, we may suppose that Pierre *is* Gibbard, a bilingual speaker of French and English, whom I will call 'PG'. We may also suppose that PG learnt both languages by the direct method, rather than one as his native language and the other by translation from it as a second language. PG has native speaker fluency in both languages. In particular, his use of both 'chien' and 'dog' is entirely normal, and neither use is parasitic on the other. He learnt both words in the usual way, through experiencing adults applying them in some cases and withholding them in others. We consider the two sentences with respect to the very same time. Thus it is straightforward that the total evidential circumstance is the same.

Suppose that PG is facing an animal that looks rather like a dog, but also rather like a wolf. He hesitates to apply the word 'chien', and hesitates to apply the word 'dog'. However, we may suppose, he happens to have experienced others applying the word 'chien' to dogs that looked slightly more like the animal now facing him than did anything to which he has experienced others applying the word 'dog'. As a result, he feels slightly more disposed to apply the word 'chien' to this animal than he is to apply the word 'dog', so he gives slightly more credence to the sentence 'Voici un chien' than he does to the sentence 'Here's a dog'.

Why doesn't PG simply observe that 'chien' and 'dog' are synonymous (for they are indeed synonymous, and he understands both), and so determine to give the corresponding sentences the

very same credence? Well, we may assume that although he is pretty confident that they have the same extension, he isn't perfectly sure; it may occur to him that 'chien' and 'dog' might conceivably differ in extension with respect to some genetically marginal animals, in virtue of differences in their use. He gives that possibility a low but non-zero credence, and refrains from believing outright that 'dog' and 'chien' are coextensive. Since his evidence for applying the term 'chien' to the animal in front of him is slightly better than his evidence for applying the term 'dog' to it, his credences differ accordingly. The point does not depend on social externalism in semantics.<sup>3</sup> Even if we treat the semantics of PG's idiolect as determined solely by his own use, he may not be certain that his own dispositions to use the words determine exactly the same extension for them. Nor is it crucial to imagine PG thinking in such metalinguistic terms. He might just think to himself, mixing languages, 'There is a slight possibility that it's a chien but not a dog'; Gibbard treats thinking as occurring in 'one's version of a public language that one speaks and understands' (p. 117).

In this story, PG does not seem irrational. Rather, he seems to be assigning credences as he ought, proportioning them to his evidence. Thus, for the same evidential circumstance, the credence that PG *ought* to assign his sentence 'Voici un chien' is slightly greater than the credence that he *ought* to assign his sentence 'Here's a dog'. Thus, by Gibbard's criterion for analytic equivalence, the two sentences are not analytically equivalent for PG. Since analytic equivalence is necessary for synonymy, the two sentences are also not synonymous for PG. This difference in meaning derives from a difference between the words 'chien' and 'dog', so they too differ in meaning, for PG. That is hardly a desirable result. By normal standards, 'chien' and 'dog' are a paradigm of synonymy across languages. Why should a trivial, idiosyncratic difference in associated experiences constitute a difference in meaning?

The argument might fail for more objective oughts, less sensitive to the agent's total epistemic circumstance. But Gibbard emphasizes that his oughts are subjective. They are 'applicable in light of information the agent has' (p. 76); 'an ideal agent must be able to recognize what they tell

her to do' (p. 118). Does that element of idealization provide an escape route from the argument? Gibbard amplifies it thus: 'These oughts are ideal in that they ignore the costs of thinking and our limited capacities to think matters through accurately' [p. 116]. He seems to intend just to idealize away from our computational limitations. But PG's problem is not a computational limitation. Merely increasing the computational power and reliability of his thought processes would not enable him to resolve his doubts as to whether 'chien' really has exactly the same extension as 'dog', even as he uses them. In particular, giving a higher credence to 'Voici un chien' than to 'Here's a dog' does not force any mathematical incoherence in PG's subjective probability distribution over sentences.

Could Gibbard insist that since 'Voici un chien' and 'Here's a dog' do in fact voice the same thought for PG, giving the sentences different credences is *ipso facto* incoherent, a lapse from perfect rationality of which no ideal agent would be guilty? That response would in effect give up Gibbard's whole project, for it invokes an independent standard for sameness of thought and therefore sameness of meaning ('the meaning of a sentence is the thought that it voices', p. 27). Since Gibbard is trying to explain synonymy in terms of what the agent's credences subjectively ought to be, he cannot also explain what the agent's credences subjectively ought to be in terms of synonymy.

Could Gibbard say that his oughts are defeasible, and that PG's doubts constitute a defeater? That would not fit Gibbard's relativization of his subjective oughts to total epistemic circumstances. In any case, he explicitly insists that his oughts are 'exceptionless oughts of rationality'. They are 'not pragmatic but purely linguistic and epistemic' (p. 114).

Yet another conceivable line is to equate synonymy with *approximate* equality in the relevant credences the agent ought to have. But that line creates unnecessary logical problems. Approximate equality is a non-transitive relation: sometimes  $x$  is approximately equal to  $y$ , and  $y$  is approximately equal to  $z$ , when  $x$  is not approximately equal to  $z$ . Since exact sameness in any given

respect is a transitive relation, exact sameness in meaning cannot be equivalent to approximate equality in the relevant credences the agent ought to have. A non-transitive standard for sameness of meaning would also make difficulties for analysing the transmission of information by testimony along a chain of several informants.

Perhaps the safest response for Gibbard is just to emphasize that on his view the slight difference in associated experiences between 'chien' and 'dog' brings only a slight difference in meaning, and to deny that such a result is anything to worry about. The trouble is that languages lose much of their point if they do not enable communication between people with different courses of experience. Such differences are just what we might hope linguistic meaning would abstract from. But since PG's word 'chien' is not synonymous with PG's word 'dog', on Gibbard's view, my word 'dog' cannot be synonymous with both his word 'chien' and his word 'dog'. Since PG's use of both French and English is normal, Gibbard's standard for inter-personal synonymy must also be ultra-sensitive to individual credences. Meanings will be far too idiosyncratic to be communicated effectively.

It gets worse. The example involves a genuine asymmetry in PG's use between 'chien' and 'dog'. He really did feel slightly more inclined to apply 'chien' than to apply 'dog' to the animal he faced. Gibbard's theory recycles that asymmetry in use as a difference in meaning. But now consider PG\*, who resembles PG except that his use of 'chien' and 'dog' is perfectly symmetrical. In any total evidential circumstance, PG\*'s credence for 'Voici un chien' and 'Here's a dog' is exactly the same. Nevertheless, I will argue, a further aspect of Gibbard's account still yields a difference in meaning between 'chien' and 'dog'.

Gibbard's full criterion for the analytic equivalence of two sentences is not just sameness of the credence one ought to have for every total evidential circumstance. He requires more: 'If two sentences are analytically equivalent, then the credences one ought to give them, for a given epistemic circumstance, must be equal under any intelligible supposition' (p. 122). Now although

PG\*'s use of 'chien' and 'dog' is in fact perfectly symmetrical, he is not perfectly certain of that fact. For all he is certain of, there may be some slight asymmetry in his use of 'chien' and 'dog' that makes them differ in extension over a few marginal cases. He does not outright believe that, as he might put it, every chien is a dog and every dog a chien. For PG\*, there is a slight epistemic possibility that some chien is not a dog and a slight epistemic possibility that some dog is not a chien, even though in any given case he gives equal credence to each possibility. For convenience, I will assume that PG\* uses the truth-functors & (conjunction) and ¬ (negation) to combine French and English sentences. Thus for PG\*, '(Voici un chien) & ¬(It's a dog)' is an intelligible supposition, as is '(It's a dog) & ¬(Voici un chien)'. Let  $Cr(X | Y)$  be PG\*'s conditional credence in X on the supposition Y. By the requirements of probabilistic coherence, equations (1) and (2) ought to hold (we may additionally assume that they do in fact hold):

$$(1) Cr(\text{Voici un chien} | (\text{Voici un chien}) \& \neg(\text{Here's a dog})) = 1$$

$$(2) Cr(\text{Here's a dog} | (\text{Voici un chien}) \& \neg(\text{Here's a dog})) = 0$$

Thus the credences that PG\* ought to give the sentences 'Voici un chien' and 'Here's a dog' under the supposition '(Voici un chien) & ¬(Here's a dog)', which is intelligible to PG\*, are maximally different in the given possible evidential circumstance. Therefore, by Gibbard's criterion, PG\*'s sentences 'Voici un chien' and 'Here's a dog' are not analytically equivalent. Since analytic equivalence is supposed to be necessary for synonymy, PG\*'s sentences are also not synonymous. As before, that comes down to a failure of synonymy between the words 'chien' and 'dog'. For the same reasons as before, various potential escape routes are closed to Gibbard. But this time there is no asymmetry in use between 'chien' and 'dog' to mitigate the result. Effectively, it is just the

phonetic difference of word and PG\*'s caution that deliver the supposed difference in meaning. Gibbard's theory individuates meanings far too finely to yield a plausible account of meaning.

Could one rescue other parts of Gibbard's theory by somehow marginalizing his account of synonymy? There is little prospect of that, because Gibbard himself gives synonymy a central role, by explaining meaning in terms of synonymy and deflation:

Pierre's sentence 'Les chiens aboient' means DOGS BARK in that it means the same as my present sentence 'Dogs bark,' and this, by deflation means DOGS BARK. (p. 179)

Gibbard's hyperbolic demands on synonymy doom such a strategy to failure. It is utterly improbable that any of my sentences meets his standard for synonymy with Pierre's sentence. The problems for his account of sameness of meaning spread through his overall theory of meaning. Nevertheless, it worth checking this conclusion from a different angle, by considering some of Gibbard's claims about inference.

## II

Gibbard explains: 'The supposition that John is a married bachelor [...] is unintelligible' (p. 122). Does such insistence withstand scrutiny? Suppose that Harzick, a native English speaker, concludes on reflection that 'Bachelors are unmarried', though plausible on first hearing, is an over-hasty generalization. He gives some positive credence to the supposition that John, who is officially still married but has been separated from his wife for twenty years, during which time he has been

leading the life of a stereotypical bachelor, is indeed a married bachelor. Harzick suspects, although he is not sure, that similarity to the stereotype may trump the usual association of 'bachelor' with 'unmarried'. I do not myself endorse that suspicion; I believe that John is married but not really a bachelor. Nevertheless, I do not find Harmick's supposition unintelligible even on *my* understanding of 'bachelor' and 'unmarried'. I can easily see how a rational, fully competent native speaker of English could come to give some positive credence to 'John is a married bachelor', without losing their competence or changing the meanings of their words. I find nothing *unintelligible* about the supposition that I myself am mistaken in believing that all bachelors are unmarried.

To change the example, here is Gibbard's paradigm of an analytic inference: 'If I accept "Snow is white" and am warranted in doing so, then I ought immediately to infer "Something is white"' (p. 116). We may assume that the argument from 'Snow is white' to 'Something is white' is in fact logically valid. Nevertheless, consider another native speaker of English, Noman. He suspects that the underlying semantic structure of 'Snow is white' may be that of a universal generalization over bits of snow. On that hypothesis, one can regiment the premise and conclusion in a first-order language as  $\forall x (\text{Snow}(x) \rightarrow \text{White}(x))$  and  $\exists x \text{White}(x)$  respectively. But that argument is invalid in first-order logic, because the premise is vacuously true and the conclusion false in any model in which the predicates 'Snow' and 'White' both have empty extensions. Noman neither stipulates nor positively believes that 'Snow is white' is a universal generalization; he just takes the hypothesis that it has the underlying semantic form of one seriously enough not to make the immediate inference from that premise alone to the conclusion. In the circumstances, it is quite implausible that Noman is changing the meanings of his words, or losing his native speaker competence, or failing to do his semantic duty. Native speaker competence is not a mode of access to the sorts of insight that professional semanticists seek to obtain, and it is of such matters that Noman is unsure.

The point is not that Gibbard has made an unfortunate choice of examples. Similar problems would arise for other examples too. There just are no semantically imposed obligations of the sort

that Gibbard postulates.<sup>4</sup> Of course, one could idealize to beings disposed to close their acceptances under genuine logical consequence (not just under what they take to be logical consequence). But that would not show being so disposed to be a semantically imposed obligation.

An *ad hominem* response on behalf of Gibbard would be to argue that the duty to proportion one's degrees of belief to probabilities on one's evidence along the lines of *Knowledge and its Limits* can make demands that similarly transcend the agent's own perspective.<sup>5</sup> In reply, several points are relevant. First, evidential probabilities in the relevant sense are not 'subjectively applicable—applicable in light of information the agent has', as Gibbard demands of his subjective oughts (p. 76). For one is not always in a position to know what one's evidence is; consequently, one is not always in a position to know what the probabilities on one's evidence are, on my view. Second, if knowledge requires outright belief, then the agent's failure to form relevant outright beliefs in the cases above implies that the agent lacks relevant knowledge, and therefore relevant evidence, on my equation of total evidence with total knowledge. For example, since PG and PG\* do not believe outright that all chiens are dogs, they do not know that all chiens are dogs.<sup>6</sup> Third, any duty to proportion one's degree of belief to probabilities on one's evidence is an epistemic matter, not a semantic one. Harzick and Noman are not violating any rules of English. Thus the analogy does not help Gibbard.

In sum, checking the inferential side of Gibbard's theory of meaning reinforces the conclusion already reached in section I through examining his account of synonymy: the tight connections between semantics and epistemology on which he relies are, as we might already have suspected, a myth.

### III

So far I have not mentioned one of the main ambitions of Gibbard's book: to provide the underpinnings in the philosophy of language for an account of normativity consistent with naturalism. Since his account of meaning is itself given in normative terms, one might suspect some kind of circularity. However, as he explains, the position is much subtler than that.

Gibbard distinguishes between properties and the concepts with which we think of those properties. We can think of the same property under quite different concepts. Gibbard's naturalism says that all properties are natural, the sort of properties we think about in doing natural science. His naturalism does not say or imply that all concepts are naturalistic, the sort of concepts we think with in doing natural science. On his view, normative concepts are not naturalistic, but that does not mean that in thinking with them we think of non-natural properties. Indeed, in thinking with non-naturalistic normative concepts, we may even be thinking of natural properties, the same properties we can also think of with naturalistic non-normative concepts. What goes for concepts also goes for meanings. Thus Gibbard's naturalism does not equate the meanings of a normative expression with the meaning of some naturalistic expression. Since his theory of meaning is couched in normative terms, it is not itself a naturalistic theory. Nevertheless, it is consistent with his naturalism, because it posits no non-natural properties. Indeed, on his view, although semantic concepts are normative and so not naturalistic, they may even be concepts of natural properties. There is no circularity in using normative terms to give a theory of meaning for normative terms, any more than there is in using semantic terms to give a theory of meaning for semantic terms.<sup>7</sup> Notoriously, the point of a theory of meaning is not to explain words to people who do not already understand them, but to move from an ordinary speaker's implicit competence with those terms to an explicit theoretical understanding of their meaning.

Gibbard calls his account of normative terms ‘expressivist’, for their distinctiveness is supposed to consist in analytic links to action. ‘The point of normative claims’, he writes, ‘is to tie in conceptually with action’ (p. 227). For example, ‘I can’t consistently believe I ought right now to leave my burning building and decide to stay. Naturalistic thoughts, in contrast, lack this conceptual tie to action’ (p. 224). Presumably, we are supposed to make sense of such talk of a ‘conceptual tie’ between thought and action by analogy with more familiar conceptual connections, such as those from ‘bachelor’ to ‘unmarried’ and from ‘Snow is white’ to ‘Something is white’. But in section II we saw grave reason to doubt that there are any such conceptual connections, at least of anything like the sort Gibbard envisages, even in those familiar cases. But if postulating Gibbard-style conceptual ties doesn’t work in the supposedly easy cases, from thought to thought, we can hardly expect it to work in the harder cases, from thought to action. There obviously *are* connections of some sort between normative thought and action, but a first step towards understanding them is to stop modelling them on a myth.

If there is no Gibbard-style conceptual tie from thought to action, what should we think of the man who believes that he ought right now to leave his burning building but decides to stay? Here is one schematic suggestion. If nothing is wrong with his belief, then he *knows* that he ought right now to leave the building, so indeed he ought right now to leave the building. In that case, since he decides to stay, he is deciding to do something that he ought not to do, so something is wrong with his decision. Therefore either something is wrong with his belief or something is wrong with his decision. Of course, we reached that conclusion by reasoning, so the reasoning had better be good. But reasoning can be good without being underpinned by Gibbard-style conceptual connections. Much ordinary human reasoning is good in the way that knowledge is good, it is knowledgeable, with no need of Gibbard-style conceptual connections.<sup>8</sup>

Of course, Gibbard’s belief in conceptual ties between normative thought and action is far from idiosyncratic. That is partly why it is so important to challenge the belief. The challenge comes

not from metaethics but from the philosophy of language. Quite generally, the evidence just does not support the hypothesis that linguistic competence embeds conceptual ties, whether between thought and action or between thought and thought. One of the main sources of resistance to a robustly realist and cognitivist conception of normativity, and indeed of morality, has been the difficulty of explaining on its terms the supposed conceptual ties from moral thought to action. If there are no such conceptual ties to explain, then one of the main sources of resistance to a robustly realist and cognitivist conception of normativity, and indeed of morality, is an illusion.<sup>9</sup>

## Notes

- 1 All page references are to Gibbard (2013) unless otherwise specified.
- 2 I have corrected an obvious slip here: Gibbard has 'Les chiens aboient' instead of 'Voici un chien', which Gibbard treats as its French equivalent (p. 119). Daniël Hoek (p.c.) points out that 'C'est un chien' would be a better choice, since unlike 'Voici un chien' it matches 'Here's a dog' in being an indicative rather than imperative sentence. That would only help my arguments, but I have followed Gibbard's choice.
- 3 Gibbard's semantics is thoroughly individualistic. Although he expresses the hope that an analogous non-individualistic account can be given, it is far from clear that it can. What is the social analogue of the epistemically transparent phenomenal qualities Gibbard assumes?
- 4 See Williamson 2007, chapter 4.
- 5 See Williamson 2000, chapter 10.
- 6 The fine-grained individuation of the contents of knowledge assumed here is anyway needed for Gibbard's project.
- 7 For Tarskian reasons, the metalanguage may have to exceed the object language in expressive resources, but that is a different issue. Although Gibbard would ultimately need to address it, that is far from the most urgent of the problems his account faces.

8 See Williamson 2009.

9 An earlier version of this material was presented to a course at the Inter-University Centre in Dubrovnik; thanks to the audience for helpful discussion.

## References

Gibbard, Allan 2012: *Meaning and Normativity*. Oxford: Oxford University Press.

Kripke, Saul 1982: *Wittgenstein on Rules and Private Language*. Oxford: Blackwell.

Quine, Willard van Orman 1951: 'Two dogmas of empiricism', *The Philosophical Review*, 60, pp. 20

-43.

Williamson, Timothy 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, Timothy 2007: *The Philosophy of Philosophy*. Oxford: Blackwell.

Williamson, Timothy 2009: 'Probability and danger' *The Amherst Lecture in Philosophy*, 4, pp. 1-35. <http://www.amherstlecture.org/williamson2009/>