

‘Boghossian’s Blind Reasoning’, Conditionalization, and Thick Concepts. A Functional Model

Olga Ramírez (Saint Louis University, Madrid)

Introduction

The present paper attempts to make a case for what I will call ‘functionally structured concepts.’ These are concepts whose inferential rules can be explained as being functionally mediated. As an example of such concepts, I will focus specifically on the case of ‘thick ethical terms’. The background of the proposal will be delivered by Boghossian’s (2003) paper on ‘Blind-Reasoning,’ because it poses the perfect context against which the considerations I will be making can be best appreciated.

One central concern of Boghossian’s ‘Blind-reasoning’ is to respond to the problem of defective concepts; these are concepts whose application rules allow us to make questionable inferences. The problem is to determine when we can legitimately be entitled to apply conceptually introduced inferences, and when we cannot. In Boghossian’s context, determining this issue is essential to his attempt to explain warrant transmission in conditional inferences, via the possession of the concept of the conditional. That is, in virtue of what are we entitled to the transition from premises to conclusion in such inferences. But the issue has relevance for itself.

Boghossian follows the lead of Russell’s (1905) treatment of non-referring expressions via existential quantification, which was later expanded by Ramsey

(1929) and Carnap (1966) to theoretical terms in science¹. Similar to the cases treated by Boghossian, the worry here was to avoid the danger of expanding our metaphysical commitments for semantic reasons. The existence and use of a concept should not mislead us into assuming the existence of a corresponding entity, or property, satisfying its cognitive content. These could sound trivial if it were not for the treacherous trap that contemporary philosophy has been holding us in: either fall back to empiricism, or accept that there is no way to prove how our concepts relate to reality and that there are no privileged discourses in this regard. This attitude has, for a good while, discouraged further semantic analysis with elucidatory existential purposes.

1. Pejoratives and Thick Ethical Terms

1.1. The challenge that defective concepts pose

Among the cases of defective concepts treated by Boghossian, is that of pejoratives. The inferential rules for the term 'boche' would legitimise us in inferring that someone is cruel just because they were German. Clearly, we cannot take for granted, on merely semantic grounds, that there is a property or disposition towards cruelty that Germans would exhibit, for the simple fact of being German. The possibility of this, and similar concepts, appeared to defy inferentialist proposals. If the meaning of a term is determined by its inferential rules alone, new concepts could be put to work by simply fixing any arbitrary inferential rules we please; and then backwards, appeal to concept possession to claim entitlement to the inferences made.

The relevance of this point, first introduced by Prior (1960) with his connective 'tonk' in the field of logic and followed by natural cases like 'boche', actually expands beyond the realm of the offensive, purposely erroneous, misleading or extravagant. Nor is it, in my view, fixed simply by appealing to coherentist requirements as Belnap (1962) proposed. Rather, it opens a more general breach in the very heart of our conceptual confidence, demanding from our concepts some non-purely-semantic legitimisation. Boghossian's (2003) position, as I see it, is precisely an attempt to acknowledge this problem.

¹ Boghossian refers to Paul Horwich's "Stimulation, Meaning and Apriority" (150-169) in Boghossian and Peacocke eds. (2000), *New Essays on the A priori*, Oxford University Press, in his treatment of this problem.

1.2. Williams, Williamson and the problem of ratification

Against the inferentialist position, Williamson (2003, 2009) has been claiming that one can understand a concept without actually being willing to infer according to the rules that the inferentialist sees the concept to be constituted by. He sees the inferentialist as being compelled to the claim that one cannot understand the concept without being ready to so infer. This is itself a disputable claim. But the decisive question here is not whether one can understand the concept without being willing to use it oneself in such a way but, rather, as Bernard Williams (1985) put it in a similar context, whether one can understand it without being willing to ratify the correctness or truth of its use by someone else. It should be noted that Williamson's point is not simply to reject that the meaning of the concept is constituted by specific inferential rules, but that it should consist in any specifiable content at all. This seems to imply that he must be ready to assert that one can grasp a concept whilst refusing to acknowledge that the truth of its application should imply any intersubjectively shared content one would be self-committed to; that is, refuse to acknowledge that it should imply any shared content we could take for granted and agree upon beyond the mere referential. If this were right, one may wonder what kind of a *factive* (Williamson 2000) statement (to use Williamson's own terminology) one is then asserting and recognising as having been asserted in such a case. But let us leave this here, for the moment, and restrict ourselves to pejoratives first.

Pejoratives certainly pose a specific case. Pejoratives are thick concepts, concepts that include an evaluation; a negative one. It is not surprising, thus, that Williamson should implicitly appeal to some of the problems posed for thick ethical terms by Bernard Williams (1985) in his metaethical writings. Williams was concerned with our attitude towards the ethical statements of foreign communities, when expressed in thick ethical concepts that we do not share. If we do understand the concepts (have been previously trained in their use) we must be ready to assent to their correct application in suitable circumstances. We should acknowledge their truth just as much as that of our own. This does not mean that we should be under any obligation to assert those statements ourselves. On the contrary, it would be natural to feel some reticence to assert concepts that we do not identify with and are not our own. Williams saw nothing awkward or wrong in this dual attitude. Similar situations are found in which we find ourselves perfectly ready to assent to the truth of statements whose selfassertion we would refuse. In this context, he uses the example of a culture where men and women are expected to use different concepts to name the same things. It is inappropriate for women to use men's concepts, and the other way round (Williamson 1985, 144-145)². Interestingly, this example is also considered by Williamson (2000) in his

² It should be noticed that the reasons why, in this specific example, we would not use the word are of quite a different sort. They do not have to do with the rejection, dislike, not even estrangement about the specific conceptual content, but are simply external reasons of politeness: "I don't get your car because it is yours and I know you would feel offended if I do, that doesn't mean I have anything against driving it," so to put it.

discussion of pejoratives with a similar purpose. In a parallel way, Williamson sees his own position as saying that one can understand a concept, be willing to ratify its correct application as true in the 3rd person case, but refuse to assert this same statement oneself.

Whether Williamson thinks, in general terms, that one can make such a difference between 3rd and 1st person assertions, or if he would just restrict it only to the case of pejoratives is not clear. However, if the claim were a general one, if Williamson thinks understanding a concept implies readiness to assent to its true application in the 3rd person case, we still have the question mentioned before, regarding the content of the commitment this person is assuming, of: what kind of shared conceptual content are they assenting to? But, leaving this point aside because it is an independent one, I do not think that Williamson would be interested in insisting on the possibility of refusing self-assertions of statements, whose truth has been accepted on semantic grounds, as a general possibility³. It is a false step, in my view, in Bernard Williams' context as applied to thick ethical concepts because Williams actually defends the cognitivist position, and he sees the evaluative aspect of thick ethical terms as part of their cognitive content. Therefore, refusing self-assertion would amount to something like: "Yes, what he says is true. Therefore, it is true that this behaviour is condemnable, meaning that it is a fact that this behaviour is condemnable; but, in ratifying this, I acquire no responsibility at all with the possible consequences of these condemnations and evaluations being taken at face value, since this is not a concept of my own."

Williams could try to avoid this problem by including some culturally relative index, something that would be more consonant with his defence of 'non-objectivity in ethics', but that is not what he does. He seems, rather, to adopt some kind of deflated factuality that could be taken to correspond to such statements. The problem, however, is strengthened for realist positions, because in their case it adopts a much more problematic ontological form.

1.3. Williamson's Proposal for Pejoratives

One may wonder what Williamson would consider us to be ratifying as true from the 3rd person perspective by thick ethical terms, and whether he would be ready to expand his defence of the dual attitude regarding 1st and 3rd person assertions to their case. Regarding pejoratives, however, Williamson's proposal does succeed in making his claim plausible. We could understand a pejorative concept and assent to

³ There are, of course, cases where we refuse to assert statements we would ratify as true. Take, for example, contexts requiring special sensitivity with someone's psychological state, like telling a person who is recovering from some deadly illness, that a loved one has been killed. However, the problem, in cases like this, affects the relevance of the asserted in a given context, and not our readiness or refusal to commit ourselves with the fact expressed by the semantic content.

a 3rd person assertion without acquiring a commitment to any conceptual content beyond the mere referential. His idea is to see the pejorative evaluative component of the concept as a conventional implicature in Grice's (1989) sense, discharged merely by its direct use.

So, while in understanding the term we come to understand that this value charged aspect will be implied whenever the concept is used, it is just by direct use that the pejorative implication is released. That is, this pejorative aspect belongs neither to the cognitive content of the term, nor to what is asserted as true when the term is applied (Williamson 2009, 24). So, apparently, we could now safely say: "Yes, what he says 'Rilke is a boche' is true," without this having any implications for myself. Since I am saying no more than, yes, it is true that Rilke is a German. The cognitive content of the concept boche reduces to its referential term, and the referential term for 'boche' is just the same as for 'German'⁴. You can judge for yourself whether this is a plausible option for pejoratives (play through, for example, Geach's (1960, 221-225) critique to see if this is compatible with giving an account of the logical use of sentences containing these terms as antecedence of conditionals, in indirect speech and so on).

This explanation of pejoratives already gives a sense of the 'referentialist' account of meaning that Williamson is heading towards in his opposition to inferentialism. He redesigns the contrast between what he considers to be a form of meaning descriptivism, which would reduce meaning to a semantic definition, and a referentialist account of meaning, that claims that it is through experience with referential instances of the concept, and appealing to such an extensional basis, that meaning should be understood. From my point of view, the contrast, so posed, is in itself questionable, because, as Gareth Evans (1973) argued, it is precisely through experience with the referent, that specifications of content in the descriptivist position take place. On the other hand, Williamson seems to assume that specifications of meaning must have some kind of static definite sense, which is not necessarily the case either.

Boghossian's proposal, as I see it, will offer a way to account for the possibility of review. I am not going to enter into lengthy discussion here because it

⁴ Imagine a parallel pejorative content for Jews, calling them "Jiws", for example. Alternatively, take an already existing concept such as the parallel case for blacks, calling them "Nigger". Imagine now, such words shouted out loud in a Neo-nazi parade at some Jews and blacks. We would have to ratify: Yes, what the Neo-nazi said when he shouted at the blacks and Jews, saying they were Niggers and Jiws is true.

One may also wonder what the difference is between shouting out something like "you blacks" or "you Jews" and developing a new word that, as opposed to the former, could not be used, but with the pejorative sense. I thank my colleague, Chris Earlam, for an insightful talk on this point.

is not the main focus of my paper, nor is, in my mind, Williamson's own sketch of his position, in the remarks made on the alluded papers, developed enough to do so. However, it will be helpful to have his referentialist orientation on the background to get a better grip of the problems we shall be discussing.

1.4. Williamson's Conventional Implicature Model and Thick Ethical Terms

One interesting question is whether Williamson's Conventional Implicature Model could be defensible for thick ethical terms, and if not, why not. In the case of pejoratives, and as an argument against a possible inferentialist reply arguing that understanding a concept is not necessarily being ready to infer according to its introduction and elimination rules, but to know that you should so infer, Williamson defends the possibility of a form of entanglement thesis.

A further problem faces even the watered-down inferentialist account of understanding pejoratives. Someone might grow up in a narrow-minded community with only pejorative words for some things, in particular with only the pejorative 'Boche' for Germans. He might understand 'Boche' as other xenophobes do without understanding 'German' or any equivalent non-pejorative term. He would be unacquainted with Boche-Introduction and any similar rule. Thus not even knowing how to infer according to Boche-Introduction is necessary for understanding 'Boche', or for having the concept that it expresses (Williamson, 2009, 11).

The idea is that we could ignore, or be unaware of, the existence of two components in the concept. Therefore, contrary to what the inferentialist claimed, to understand the concept is not that we must be ready to infer from the one to the other, but we do not even have to be able to distinguish two components. Here again, Williamson plays up the similarities between pejoratives and questions elicited by standard thick concepts. It is doubtful, though, whether this possibility is at all compatible with his proposal for the understanding of pejoratives in terms of conventional implicatures. That is, whether the language users need not be able to sort out the referential or application conditions of the term, even if they lacked the term 'German,' from its offensive implications.

If we consider the process of conceptual acquisition of such a concept, we may imagine that the apprentices will first have learnt to sort out a type (coincident with Germans, even if they do not know this) according to the non-pejorative criteria (as the offensive implications are, per definition in Williamson's model, not needed for that). But this would not have been enough. The instructor must then make clear something else, something extra: "have you recognised those I mean?" he would ask. "Well, there is something wrong about them..." (or maybe, closer to Williamson's understanding: "well, we dislike them and want to insult them saying that they are cruel") he may add. He may say, for example: "look, those are the

boches and they are cruel and distasteful and so we say". But the separation (before and after the 'and') is anyway clearly made.

Williamson could insist that they learn it altogether, as referred to each instance in the learning process. But what they have to fix their attention on, in order to be able to apply the term, would be the referential aspect, according to Williamson's own definition. Some apprentice could have later on argued without contradiction: "yes, I know those you mean with 'boches', but there is nothing wrong with them, they are neither cruel nor disgusting at all." Or maybe: "yes, I know who you mean with the 'boches', but I don't want to insult them, they are perfectly alright." This speaker will have to be taught that he cannot say that because if they are the one (the specific type in question) then (with an enough amount of 'Bildung' the trainer might want to be claiming) he will 'see' or he can take for granted that they are the other (cruel).

Just because we make this *extra step* we can call them 'boche,' if he should not be ready to do so, then the type should not be called 'boche;' or maybe, more according to Williamson's conventional implicature reading, he may say that the apprentice cannot then use the word. He should not use it if he does not want to insult such people, because the offense belongs to what we want to say with it. We want to make *them* (that group you have learned to sort out) angry, we want to say that because we dislike them and we know that this historical association with their cruelty will offend them. Actually, the reason why people, in Williamson's Conventional Implicature Model, might not want to apply the term in the 1st person is precisely because they do not want to insult such people; that is, because they distinguish between the group of people and the insult.

In the case of thick concepts, the disentangling problem was supposed to be that we would not be able to determine the extension of the concept, in the first place, without the aid of the evaluative component, and without taking into account evaluative considerations. According to McDowell (1981), we could not even sort out any non-evaluative patterns, once the concept is constituted, making true their application.⁵ In Williamson's proposal, I have just argued that, unlike what he is suggesting, this must be granted, because it is the non-pejorative element that determines the truth conditional meaning of the term. The very idea of the pejorative component being a conventional implicature, speakers could refuse to 'discharge' by declining 1st person use, while acknowledging the truth of 3rd person applications,

⁵ Remarkably, nothing like having to identify oneself with the sensibility of the language users to fully understand the sense of the term (as is sometimes argued in the moral context) would seem appropriate here. Although, maybe we should not underestimate what intensive, and sufficiently penetrating, training is capable of achieving.

implies their capacity to separate both. Surely, the difference between the simple thick case and the pejorative would have to be taken into account, because, in the pejorative case, the evaluative aspect is itself a thick concept. But, for the purpose of questioning, whether a distinction can be traced between the application conditions and what the concept beyond it says, the problem posed is similar enough.

Why could it not be claimed that, as in Williamson's pejorative case, the fact that there is no non-evaluative concept distinguishable by the users does not mean that we do not have to learn to sort out some application conditions on whose behalf the moral evaluation and thick concept are applied? Expressivists have famously argued this way, but failed because they situate the application conditions at a naturalistic non-conceptual level. At a purely physical level, the members of the extension of a thick ethical term may have nothing in common, as McDowell claimed. But, as I have argued elsewhere (Ramírez 2011, 106), the application conditions, just as it happens in Williamson's example, need not be so basic. The term 'German' does not have as its application conditions a natural kind either, there are nonetheless application conditions⁶ for its use, and they are different from what is inferred on their basis in the composite concept boche.

The application conditions for thick ethical terms, those aspects we must prove before applying the term (even if we do not have a special term for them), can be conceptually specifiable in a similar fashion and distinguishable from the evaluation made on their behalf. Once this is granted, why should a model such as Williamson's not be adequate for thick ethical terms? The application conditions for the term, those we must be careful in proving, would pose the 'referential term' in Williamson's vocabulary. The evaluation, if we proceed with the application of Williamson's model, would be explained as a conventional implicature discharged in use, but not belonging to the concept's cognitive content.

What benefits do we obtain through this account? Well, first, to solve the problem of having to ratify the right application of foreign thick ethical concepts we do not share without acquiring any responsibility with the truth of the asserted (now, 'conventionally implied') evaluations. In assenting to the truth of a 3rd person statement, we would simply confirm that the application conditions of the thick term are satisfied; declining 1st person use, we decline to release the moral evaluation in implicature form. Second, it would serve the anti-inferentialist claim that we can

⁶ The fact that we do not always make the effort to take everything into account in order to determine whether someone is a German or not, does not mean that we do not recognize them as relevant to determining whether someone is German or not. Discovering that some aspects were not fulfilled would make us revoke our original judgment. This also allows, of course, some vague margins, accepting specialist's judgments and redefinitions according to law or whatever.

understand the concept without being ready to infer according to its putative inferential rules. But, as in Williamson's pejorative case and for the same reasons, the idea that we should be able to say still that we can understand the concept, without even making sense of the difference between application conditions and evaluative implicature, would not make sense in this model.

Why should this not be an appropriate model for the ethical case? The problem is that even if we succeed in the aforementioned kind of disentangled explanation, and so much seems quite plausible to me, there is a fundamental difference with pejoratives. Here it makes sense to ask: Why should we sort out such 'referential' characteristics for the moral evaluation at all? What is the point of sorting out such a class (based on a set of criteria whose fulfilment we have to prove) without reference to some evaluative interests of ours? This question, and not the impossibility to tell apart application conditions and moral evaluation (and both are not to be conflated) still makes sense from a non-realist point of view (Ramirez 2011, 106-107). In the case of pejoratives, there need be no special objective reasons that anyone is expected to make sense of. The connection might be quite arbitrary, expressing simply subjective disgust or a conscious decision to harm a given group of people by way of applying them a derogative value, as in the case of *boche*. In the moral case, on the contrary, I think we do tend to demand, or at least expect, there to be such a connection.

If we accept and demand rational explanations (answers to the *why* question) then we adopt a cognitivist reading. We think it can be justified why the behaviour is considered morally good, and, correspondingly, why an attribution of rightness or truth to a corresponding statement includes, as part of the justified cognitive content, the moral value too. Expressivist or projectivist approaches, on the contrary, share with Williamson's implicature model the exclusion of moral evaluations from the cognitive truth-evaluable content of the term. Moreover, we should question how big the difference is between turning the evaluation into a kind of expressive illocutionary force and turning it into a conventional implicature (that is, into some wider, but not truth apt, part of its meaning). In both cases, the evaluation is emitted just by direct use. In both cases, we know in advance that by asserting the concept we are expressing the evaluation. However, even though, at the semantic level of analysis, there are more similarities than there might at first sight seem, there is still a remarkable difference between both approaches; at least there is between the pejorative case, as treated by Williamson, and the projectivist ethical case.

Most projectivists (Blackburn 1981) do attempt to give some kind of explanation, in naturalist terms, of the relationship between 'evaluated' and 'evaluation.' This is not an arbitrary casual decision, like the decision to harm

whether or not it is deserved, in the case of pejoratives, but is supposed to be a causally explainable relationship. We are certainly not talking in terms of rational justification, but of a genealogical causal explanation: how we came, and how we were caused to attribute some value to certain behaviour. Whether or not this is right, or whether what is naturally originated is also right from a rational point of view, is not something we are judging in this model.

Someone might argue, that the introducers of pejoratives may equally claim that it is because this special group of people trigger in them such disgusting emotions that they project those emotions onto them. But, still, a projectivist model has the aspiration to be a genealogical explanation of the origins of these evaluative thick ethical terms, departing from a naturalistic picture of the world. While those first introducers of pejoratives could hardly claim their views are not theory-laden (actually, charged by deformed interpretations). Something that is not compatible with the projectivist causal model. This problem, by the way, pace projectivists own self interpretation, which does not seem to register it, affects their own account of thick ethical terms. The dilemma that projectivists are trapped in is the following: if they keep to the picture of reactions to a conceptually independent world, then they become the target of McDowell's 'Rule-Following Critique.' According to this critique, the members of the extension of a thick ethical term may have no shared features at a physical level, so projectivists could not explain what it is that guides them in applying the concept. It could not be our reactions alone, because these are shared by all thick concepts: it is acceptance ("hurray!") by all positive concepts and rejection ("boo!") by all negative ones. So, the disentangling, and the possibility of following a rule on a naturalistic basis, fails.

If, on the other hand, and in order to solve this problem, they should propose to situate the common features that guide conceptual application at a higher order conceptual level (Ramírez 2011, 106), they must give up the causal model. If we react because we previously understood something to be wrong (on rational reasons), then we are not any more in a naturalistic framework, but presupposing cognitivist explanations. However, we could say that, at least at the level of their aspirations, there is a difference between the implicature model and the projectivist one. While the second attempts to give some commonly acceptable explanation of the relationship between 'evaluation' and 'evaluated,' this is not necessarily the case by the first. The Implicature Model must admit, when not the subjective and arbitrary, at least the limitedly shared character of the connection the concept establishes. The purpose of this experimental attempt to apply Williamson's Conventional Implicature Model to thick ethical terms, was to make clear where the insufficiencies would lie. Whatever its advantages in escaping the ratification problem, it would fail

to explain the connection between the evaluated and the moral evaluation; but we do want to demand an explanation of such a connection in the moral case. On the other hand, however, the comparison was guided by the possibility of making plausible, a disentangled account of these concepts; where the shared application conditions would not be conceived in naturalistic terms, and, in this way, avoiding the fundamental problems that expressivist and projectivist disentangled accounts face.

2. Boghossian's Proposal

2.1. Advantages of an understanding of thick ethical terms, in line with Boghossian's proposal

Boghossian's position shows the possibility of a more adequate treatment of thick ethical terms, and, in my opinion, lies closer to what could be a proper understanding of pejoratives. Critiques of the inferentialist position have used the case of defective concepts to question the appeal to concept possession as a way to justify inferential entitlement. This is described by Boghossian's (2003) MEC (Meaning, Entitlement connection): "any inferential transitions built into the possession conditions for a concept are *eo ipso* entitling" (cf., 17). From Boghossian's perspective, however, the problem is not that we should appeal to semantic criteria in these terms, but that there should be no epistemic justification for the corresponding conceptually introduced inferences.

In the case of pejoratives, Boghossian does not deny that they should be considered to be genuine concepts, and actually, they are normally used as such by language speakers. The term 'Boche' is defective not for that reason but because we lack a justification for the connection between the 'referential term,' in Williamson's sense, and the pejorative evaluation. Williamson may argue that offensive concepts do not need it, but according to Boghossian, they remain fraudulent concepts for precisely that reason. The consequence of this realisation, though, or so one would tend to think, should be to rule these concepts out. But this, as Williamson argues, speaks against our sense of them fulfilling some kind of social function, even if it should somehow be a wicked one.

A possible reply, consistent with the inferentialist account, may be to say that with such concepts we actually pretend that there is a connection when there is not (we may have some capricious arbitrary reasons or subjective feelings, but no objective epistemic justification), and that is exactly how they work. That is how the offense is successful. Just as lies work by pretending non-truths to be true. There is something parasitic about such conceptual uses. The offended feels something is

said about him and his peers ‘as if it were true’. So, maybe a more plausible account of pejoratives could go along the lines of a ‘pretence’ or ‘make-belief,’ sharing some similarities with fictional uses of other concepts. Strictly speaking, however, such inferences would be wrong. I think this is a path open for the pejorative account and an analysis of pejoratives along these lines would not put into question the Boghossian proposal.

Before we see what Boghossian proposes to deal with the possibility of defective concepts, it is worth taking a look at some further examples that he introduces to explain why mere appeal to the MEC is not sufficient warranty for the epistemic validity of our inferences. The cases he considers are the hypothetical terms ‘aqua’ and ‘flurg’ (Boghossian 2003, 243–244). The introduction and elimination rules for the term ‘aqua’ would allow us to infer from x being water to its being H_2O . That is, (IA) x is ‘aqua’/ x is water; (EA) x is water/ x is H_2O . It is assumed that it is in virtue of it being (having this new property) *aqua* that water is H_2O .

The term ‘flurg’ is similarly introduced to enable us to infer that elliptical equations can be correlated with modular forms ((IF) x is ‘flurg’/ x is an elliptical equation; (EF) x is an elliptical equation/ x can be correlated with a modular form). It is because of elliptical equations being *flurg* that they can be so correlated. This way the term ‘flurg’ would allow us to infer on a purely semantic basis what Wiles (1995) needed to prove to show that the Taniyama-Shimura conjecture was right (Boghossian 2003, 244): that elliptical equations can be correlated with modular forms.

Using Dummett’s (1973) terminology we should say that these concepts, like the concept *boche*, are non-harmonic. They allow us to derive more information per elimination than the introduction rule brought in. The inferences they introduce are also non-conservative in Belnap’s sense. They make inferential transitions possible that are incompatible with those allowed in our pre-existent linguistic system, giving rise to problems of inconsistency. For Boghossian, though, the problem is not exactly that these concepts should be nonconservative. It is, rather, that the modifications and enhancements in our linguistic system that they introduce are not epistemically justified, while the concept allows you to take for granted, as an a priori matter, that they are. Some inferences may initially be inconsistent with our pre-existent linguistic system, and instead of abandoning them we may opt to introduce new, but richer adjustments in the system. This could never be accepted nor explained on a purely coherentist basis. Appealing to conservatism and consistency may, thus, be necessary, but not be sufficient.

The possibility of fraudulent concepts such as *aqua*, *flurg* or *boche* (even if in this last case its use could have a pragmatic purpose) threatens a too naïf appeal to meaning rules to justify the truth preserving character of our inferences. So the challenge that Boghossian sees is to determine in which cases we are safe to do so and in which cases not. Boghossian's response to this challenge depends on his distinction between conditionalised and unconditionalised concepts (Boghossian 2003, 245-246). Conditionalised concepts are those whose use should be made dependent upon an antecedent proof of the existential claims made by its conceptual content. Unconditionalised concepts, on the contrary, would not demand any such justification. The problem, according to Boghossian, is to use some concepts as if they were unconditionalised ones, while they are not, leaving no possibility open to prove their epistemic cogency.

Through this distinction, Boghossian seems to be suggesting some kind of epistemic foundationalism. The idea appears to be that some concepts are basic enough not to require any further epistemic justification, while the complexity of others demands us to leave open the possibility to reassure ourselves of their epistemic soundness. This appeal to conditionalisation poses a step towards restoring the notion of conceptual analysis in the classical russellian sense. Boghossian appeals more specifically to Carnap's development of Ramsey's application of the russellian quantification procedure to theoretical terms in science. Theoretical terms make presuppositions that are not directly verifiable by the observations on which basis they are applied. To prove whether the direct application of a theoretical concept such as 'neutrino' is legitimate, Carnap (1966) proposed the following analysis:

T (Neutrino)
(S) Ex (Tx)
(M) Ex (Tx) \rightarrow T(Neutrino) (cf., 245)

The idea is that when speakers use the term 'neutrino' they would be conditionalising upon the existence of something satisfying the theoretical assumptions of the neutrino theory. That is what is expressed by sentence (M). The truth of their application depends, then, on whether the sentence (S) is satisfied: Whether there are any items satisfying the theoretical assumptions made by the term in our scientific theories.

The existential presuppositions made by our concepts could in this way be subject to review. But why should this be necessary at all? In using the concept precisely, are we not identifying the existence of such individuals? This may be our immediate reaction, but the point requires careful consideration. The unconditionalised use of 'neutrino,' for example, implies that when scientists use the

term they would be applying it to some observational event, declaring it to instantiate a token of a type whose existence we are simply taking for granted. As the conditions of application are fulfilled, no further proof is required. Similarly, with the unconditionalised use of cases such as 'aqua' or 'flurg' a comparable problem would be given rise to. Whether the type in question really exists, however, or whether there are instances of such, is what we are trying to prove with (S).

Now, speaking of angels or unicorns, in the context of the Russellian discussion about nonexistents, we are presupposing that there is a referent for our terms when there is none. We actually have no conditions of application at all. Similar to the case of the neutrino, however, it could be argued that there are in fact some observations in the case of angels; some feelings, effects, or whatever, that are symptomatic of the presence of an angel as their hidden cause. I am no expert in angel theory, so I do not know exactly how this is determined, but other cases of mythological, fictional or sacred figures can be said to be simply postulated, without there being conditions of application for them at all.

The cases of non-harmonic concepts we were considering before, cases such as 'boche,' 'aqua' or 'flurg,' however, are not like this; at least not like these last ones. They do have application conditions, they do have a referential term, and that is precisely the problem. Since we consider them to be rightly applied on such a basis, their possible defectiveness can remain unnoticed. We do not, so easily, question the existence of the type sorted out, because of the sort of immediate reasoning considered before. We already seem to have sorted out a set of items that would constitute the extension of the term in question. We do not simply presuppose existence in the sense of 'Plato's beard', but we seem to confirm existence on an application basis.

To lean towards the connection of these questions with the ones regarding thick concepts and pejoratives with which we started, let us go a bit more explicitly into the problems that the non-harmonic concepts seem to pose. We saw that, according to Williamson, the terms 'boche' and 'German' would have the same referential basis. Nevertheless, the term 'boche,' considering its 'wide meaning', says more than 'German.' Following Williamson's Conventional Implicature Model, this is no problem because truth is applied to the cognitive content of the term, and this is exhausted by its referential meaning. That is, the sentences 'your boyfriend is a boche' and 'your boyfriend is a German' would have the same truth conditions, no matter if the first, when directly asserted, discharges the negative pejorative connotations. I would have to acknowledge the truth of both, as applied to my German boyfriend.

When considering the term 'aqua', once more two terms, 'aqua' and 'water', would have the same referential basis in Williamson's terminology. They both have the same truth conditions, they are rightly applied to instances of water, but the sentences: a) 'there is water in the glass,' and b) 'there is aqua in the glass' are semantically different. Sentence b) says more than Sentence a). Here, however, we cannot do away with the exceeding semantic content by claiming that it is some kind of conventional implicature. The concept 'aqua' is making a different assertion about the world's chemical composition, not just expressing some dismissal or any other kind of attitude towards the world. So, when we assert the truth of Sentence b) on the basis of an instance of water in a glass, then truth applies to the whole cognitive content of the sentence. We are saying that it is true that there is aqua in the glass, and, therefore, that there is this further chemical element there turning water into H₂O.

A similar case could be brought up for the case of flurg. The problem is, therefore, that the application conditions on which basis sentences of the type of Sentence b), constituted by non-harmonic concepts, are considered to be true, fall short in proving the rightness of the asserted content. So, what conclusions should we derive from here? Williamson's referentialist account seems to be suggesting that we focus on extensional aspects of meaning. Learning to identify the right application conditions, together with some indefinite sense of what the concept actually says acquired through social training, is all there is to conceptual understanding. However, this does not seem sufficient to protect us from the possibility of theerroneous assumptions that the unconditionalised use of these concepts compels us take for warranted.

The mere knowledge of the application conditions of a concept does not always seem to justify the truth of our assertions. This reflects upon the problem of ratification, considered at the outset, in the context of Bernard Williams' discussion of thick ethical terms. It was because we assumed that the right application unquestionably warranted the truth that we were caught in the conflict of having to ratify as true thick ethical concepts we were reluctant to use and share, and not for reasons of mere politeness. If we recollect our discussion in the first sections, the problem of applying Williamson's model to thick ethical terms was not that of disentangling the application conditions from the evaluative implications. I argued that, as long as we do not want to understand such application conditions in a naturalistic sense, McDowell's criticism of disentangling does not stay in the way. An inferentialist explanation of thick ethical terms would be a disentangled account that separates, in a non-naturalistic way, some conceptually specifiable application

conditions from their moral evaluation; when considered in this way, thick ethical concepts may be another example of non-harmonic concepts.

When the application conditions of thick ethical concepts are fulfilled, we infer that the moral evaluation applies too. However, for the reasons given when discussing the inconveniences of the projectivist and conventional implicature's approach, this inferentialist position would be a cognitivist one. Therefore, in asserting that the application conditions are satisfied we would be stating as true that the exceeding semantic content of the term, the moral evaluation, is also truthfully applied. Now, if we adopt Boghossian's proposal of conditionalisation for thick ethical terms, the sceptic foreigner, in Bernard Williams' context, now has a chance to refuse ratification requiring an epistemic proof that justifies the validity of the inference the concept takes for granted. He can admit that the application conditions are fulfilled, but, nonetheless, he will also have a chance to ask why some specific behaviour is considered morally wicked, for example.

Some may wonder why a harmonic reading of thick ethical terms may not be possible. The non-harmonic reading of thick ethical terms would take a general form, such as this:

INTRODUCTION RULE

x is thick ethical $cpt_{(1,2,3...n)}$

x has application $cond_{(1,2,3...n)}$

ELIMINATION RULE

x has application $cond_{(1,2,3...n)}$

x is (morally) wrong/ right

cpt = concept and $cond$ = conditions

While the application conditions would be specified for each particular thick ethical term (1,2,3...n), as those specific conditions (1,2,3...n) that we must pay attention to in order to determine if we have to do with an instance of application of the concept, the moral value attributed on such a basis would, on the contrary, adopt a bivalent character (either positive or negative).

If we want to offer a harmonic reading, we could propose a conjunctive (see Williamson's discussion, cf. 260) reading of the term:

INTRODUCTION RULE

x is thick ethical $cpt_{(n)}$

x has application $cond_{(n)}$ & x is morally w/r

ELIMINATION RULE

x has application $cond_{(n)}$ & x is morally w/r

x is thick ethical cpt_n

w/r = wrong/right

If we put it in this way, we would also need an independent way of proving, for each instance of the application conditions, whether it has a positive or negative moral value. The first problem of such a reconstruction is specifying what must be proved, in order to determine whether the behaviour fulfils the moral requirement. The second problem is how practical it would be to require that each child and concept user should stop to consider, in each particular occasion, whether it is fulfilled. However, the real question is: What is the purpose of having a 'thick concept' if we are to do all that proving in each case over and over again? Are we, as users of the thick concept, supposed to consider it proof dependent that some instance of a behaviour that fulfils the application conditions¹ (those specific to a given thick concept) for example, could fail to be morally blameworthy? Is it at all compatible that application conditions_i could happen without blameworthiness applying?

The purpose of having a thick concept is to assure us that if the application conditions are the case, the moral evaluation will on each occasion apply. Therefore, requiring an independent proof on each occasion does not seem to make sense. To

propose a conditionalised reading, on the contrary, does make sense, because even if the idea of having such a concept is to tie application conditions and evaluation in a generic sense, this does not necessarily mean that the alliance is an irremovable one. The risk of giving a realist reading of these concepts is to take for granted that because we once had reasons to think that some given type of behavior was morally blameworthy this evaluation belongs irremovably and intrinsically to it, as some kind of ontological fact. Boghossian's conditionalised reading offers the opportunity to prove whether the alliance we once made was a sound one that we are still ready to support.

2.2. Functionally Structured Concepts

The question is now, to what extent Boghossian's specific account of conditionalisation is delivering the tools we need to prove the reliability of those inferences that thick ethical concepts introduce. Taking as a model Russell's analysis of non-existents and Carnap cum Ramsey's proposal for theoretical terms in science has shaped Boghossian's proposal of conditionalisation and how he sees its application to the examples chosen. The epistemic legitimacy of a concept is made dependent on a quantified existential sentence, whose satisfaction we can prove. We turn the substantive 'unicorn' into a definite description and prove whether something satisfies it. The model fits theoretical scientific terms like 'neutrino' too, since the question posed is whether there really is something causing the effects taken to be symptomatic of a neutrino's presence. Similarly, when Boghossian offers an analysis of the term 'flurg,' he follows the same pattern: "If there is a property which is such that any elliptical equation has it, and if something has it, then it can be correlated with a modular form, then if x has that property, x is flurg" (p. 247).

The purpose of the concept *flurg*, if we follow the introduction and elimination rules Boghossian offers for it ((IF) x is flurg/ x is an elliptical equation; (EF) x is an elliptical equation/ x can be correlated with a modular form) is to entitle us to the transition from x being an elliptical equation to its correlation to a modular form. In the above definition, though, Boghossian offers a reconstruction of the concept *flurg* according to which what we are to prove is the existence of a new property, call it Z , that could be made responsible for the transition. The question is: why was this necessary? Clearly, this makes the parallel with the concept of the neutrino more effective, but not all concepts can be proven this way. The fact that there should be for some elliptical equations a (casually) correlative modular form, making the description extensionally true, would not provide what we need either. What is required is to show that as *a general rule* and as *a matter of necessity* all elliptical equations can be transformed into a correlative modular form; the question

is whether we can assume that they can be so transformed. This seems to require not a further entity, about which we may wonder once again how it itself makes the correlation with modular forms possible, but a proof, Wiles proof.

The case of thick ethical concepts seems to me to be similar in this sense. A conditionalised reading of a thick ethical terms cannot take the form of a quantified existential sentence, according to which we would be looking for some behaviour that makes the description true. For instance, being of the type of behaviour that is characterised by the application conditions and is, therefore, morally blameworthy. What we need is proof showing that the described characteristics of a (possible) behaviour require (deserve or should be given) a specific negative or positive moral evaluation. It is important to notice that this is not an existential requirement in Boghossian's sense, because the searched connection is to hold, even if there is no individual satisfying the description. We could imagine a behaviour that has not yet taken place or that may not ever exist, but claim that if someone were to behave in that way his behaviour would be morally condemnable.

Accordingly, we may distinguish two different ways to prove conceptual legitimacy through conditionalisation:

- a) Quantifying over entities: cases where the issue is to prove if there is something making true the theoretical assumptions that have been made about some presumed existent entities.

- b) Quantifying over functions: cases where the problem lies in proving the rightness of the connection between different entities, or, for example, between some specific behaviour and its evaluation that some concepts allow us to take for granted.⁷

⁷ This proposal was made in a talk given in Barcelona at the conference of the SEFA (2011), 4-6 April, in honour of Crispin Wright; although the analysis was not applied to thick ethical terms. However, in Ramirez (2011) and also in two previous talks, in Barcelona's SEFA (2007) and at the 31th Wittgenstein's Symposium in Kirchberg (2008), an independent proposal was made to see thick ethical terms as functionally structured. It was after encountering Boghossian's work later, that it occurred to me to put together both my proposal and his, and see thick ethical concepts in his frame with the addition of functional analysis.

I want to add that the first time I presented a first reconstruction of thick concepts on similar lines (though not completely in its actual form) in the Seminar TEC (2006) in Granada my colleagues María José Frapolli and J.Luis Liñan suggested that I should consider it from an inferentialist perspective too.

A further question is why such proof should be needed for some inferential transitions of this kind, and should we always demand it? It seems to me that proof should be required in cases where doubt makes sense and the transition is not self-evident. That is, where it is not self-evident to people who have not been previously trained to consider it so.

When we, earlier, considered the insufficiencies of Williamson's model and the projectivist one as accounts of thick ethical terms, I argued that they failed to provide a reason to explain why a given behaviour should be evaluated in a given way. As McDowell puts it, there may not be independent reasons to fix attention on such behaviour if it were not from the perspective of the moral evaluation. This, I think, is to some extent right, but need not speak for an entangled account, or a realist position. The idea, to put it in terms of our earlier discussion, is that we do need to explain why the evaluated is selected as suitable for the moral evaluation. It could well be that it is having some previous idea of what we search for in terms of 'morality' (what we consider to make behaviours *morally good*) which is what fixes our attention in selecting given patterns of behaviour as those that satisfy or go against our criteria. I have previously argued for this in Ramírez (2011).

When we talk about moral behaviour, we talk about something specific. When we say something is 'morally' good, we are not simply saying it is good in the way that a child says that getting what he wants is good, or in the way that eating a pineapple cake is good (if it is). Things can be good from different perspectives; it favours my group, or it goes against my enemy, for instance. When we consider thick ethical terms, however, when we call them so, and say they include an evaluation; we mean a *moral one*. So determining why some behaviour is 'morally good or bad' must require determining this specific sense according to which something is *morally* good/bad. I think there may be broad enough consensus in saying, in a minimalistic way, that morality deals with those relations of men, with each other and with their environment, which we want to expect from all. So, any behaviour would be morally good if *in this sense* it is good.

Different philosophers come to different and more specific conclusions regarding what would be good in this sense. Kantians, I suggest, would say that behaviours fit the bill if their general observance would equally protect the preferences (Hare, 1963), or the needs and interests (Habermas, 1983) of those affected. However you may put it, there are ways to deliver the standard, relative to whose fulfilment, behaviours are to be evaluated in a moral sense. To allow these different possibilities, I proposed (2011), in a more abstract sense, that some behaviour is morally good if it fulfils a well defined moral function. If we make this our goal, if we want to select behaviours from this perspective, then those who

satisfy this function will be the ones to be called *morally good*, and those going against it will be *morally blameworthy*.

I have left aside the problem of distinguishing between 'ethical' (in a narrow sense) and 'moral' values, as some authors do. This could well be expressed in terms of the reach width that some of these moral expectations have; their circumscription to a given community and the life of its individuals. Clearly, to pretend to talk in terms of truth when dealing with thick ethical terms whose reach is circumscribed to a given society, makes things much more problematic. Here, ratifying their use as true on the basis of the right application conditions is less advisable than in the more universally meant cases. However, if the language users, as people like Putnam have claimed, do want to make universal claims through these concepts, then adopting a conditionalised position, as Boghossian proposed, and clearly determining the standard used will be most useful. It will, nonetheless, depend upon whether we are ready to accept the used standard and behaviours fulfilling it as morally good in a shared sense, that we will accept or refuse the legitimacy of the concept. If some cultures want to claim that some behaviour will be right in a restricted sense, according to some specific conditions of their society, then truth should be made relative to those conditions.

According to these considerations, I would conclude that a conditionalised reading of thick ethical concepts is desirable, and that the quantified sentence needed would be a functional one. If there is a function such that when applied to a given behaviour it allows us the transition to the evaluation introduced by the thick ethical concept, then the concept is a legitimate one. The possibility of review to see whether new considerations change the conclusions of our proof remains, in a conditionalised account, nonetheless possible. To finish with, I just want to mention that even if I believe that cases of the type of flurg (and flurg if it were to exist) and thick ethical concepts are good candidates for the kind of functional quantification I have been proposing, there is a difference between them. It seems to me that in the case of flurg, the problem is a matter of finding a proof that shows that there is such a connection as the Taniyama Shimura Conjecture suggested. In the case of thick ethical concepts, it is not a matter of showing through proof that there is a connection, but of introducing a new standard, that is not already there, relative to whose fulfilment the evaluation is made. So here, especially, the inferential relation needs some external mediation and I have previously, informally, reconstructed how this might be represented (Ramirez 2011, 109). How exactly it is to be expressed formally is something that I will leave for now.

Literature

- Belnap, N. D. 1962. Tonk, Plonk and Plink. *Analysis* 22, 6. June: 130–134.
- Blackburn, S. 1981. Reply, Rule-Following and Moral Realism, in: *Wittgenstein: To Follow a Rule*, S. Holtzman & C. Leich (Eds.). London: Routledge, 163–187.
- Boghossian, P. 2003. Blind Reasoning. *Proceedings of the Aristotelian Society*, Supplementary Volume 77: 225–248.
- Carnap, R. 1966. *Philosophical Foundations of Physics: An Introduction to the Philosophy of Science*. New York: Basic Books.
- Dummett, M. 1973. *Frege: Philosophy of Language*, London: Duckworth.
- Evans, G. 1973. The Causal Theory of Names. *Proceedings of the Aristotelian Society*, Supplementary Volume 47: 187–208.
- Geach, P. 1960. Ascriptivism. *The Philosophical Review* LXIX: 221–225.
- Grice, P. 1989. *Studies in the Way of Words*. Harvard University Press.
- Horwich, P. 2000. Stimulation, Meaning and Apriority, in P. Boghossian & Ch. Peacocke (Eds.), *New Essays on the A Priori*. Oxford University Press, 150–169.
- Lewis, D. K. 1997. Finkish Dispositions. *Philosophical Quarterly* 47: 143–58.
- McDowell, J. 1981. Non-Cognitivism and Rule Following, in S. Holtzman and C. Leich (Eds.), *Wittgenstein: To Follow a Rule*. London: Routledge, 141–162.
- Prior, A. 1960. The Runabout Inference Ticket. *Analysis* 21: 38–9.
- Ramírez, O. 2011. Between Non-Cognitivism and Realism in Ethics: A Three-fold Model. *Prolegomena* 10(1): 101–212.
- Ramsey, F. P. 1929. "Theories," republished in his *Foundations*. In D. H. Mellor (Ed.), *Essays in Philosophy, Logic, Mathematics and Economics*. London: Routledge & Kegan Paul.
- Russell, B. 1905. On Denoting. *Mind New Series* 14, N^o 56: 479–493.
- Williams, B. 1985. *Ethics and the Limits of Philosophy*. Harvard University Press.
- Williamson, T. 2000. *Knowledge and its Limits*, Oxford University Press.
- Williamson, T. 2003. Blind Reasoning. *Proceedings of the Aristotelian Society*, Supplementary Volume 77: 249–293.
- Williamson, T. 2009. Reference, Inference and the Semantics of Pejoratives, in J. Almog, P. Leonardi (Eds.), *The Philosophy of David Kaplan*. Oxford University Press, 137–158.

Olga Ramírez (Saint Louis University, Madrid)

**'Boghossian's Blind Reasoning,' Conditionalization, and Thick Concepts.
A Functional Model**

Abstract. Boghossian's (2003) proposal to conditionalize concepts as a way to secure their legitimacy in disputable cases applies well, not just to pejoratives – on whose account Boghossian first proposed it – but also to thick ethical concepts. It actually has important advantages when dealing with some worries raised by the application of thick ethical terms, and the truth and facticity of corresponding statements. In this paper, I will try to show, however, that thick ethical concepts present a specific case, whose analysis requires a somewhat different reconstruction from that which Boghossian offers. A proper account of thick ethical concepts should be able to explain how 'evaluated' and 'evaluation' are connected.

Keywords. Meta-ethics, cognitive aspects of morality.

Citations. Reference this paper as: Ramírez, O. (2012). 'Boghossian's Blind Reasoning,' Conditionalization, and Thick Concepts. A Functional Model. *Ethics in Progress Quarterly*, Volume 3, Issue 1, pp. 31–52, available online at ethicsinprogress.org.