

# Moral Relativism

MAX KÖLBEL

There is a general philosophical problem with the status of non-objective judgements, such as evaluative or aesthetic judgements. On the one hand, even judgements in these areas seem to answer to some standard of correctness. They are not arbitrary. We often call them “true” or “false”. One can err in making them. We argue about them. On the other hand, there are reasons to say these judgements do not answer to, or describe, anything in objective reality. For if there were, a huge epistemological puzzle would arise as to how we access these evaluative aspects of the world. The general problem is to make room for non-arbitrariness, correctness and error, truth and falsity without positing epistemologically dubious objects of knowledge, such as non-natural facts etc. Let’s call this the general problem of non-objectivity. One answer to it is to say that these apparently non-objective judgements do, after all, answer to an absolute standard of correctness, and the challenge is then to come up with an epistemologically plausible story about that standard. Another answer is to insist that these judgements do not, after all, answer to an absolute standard of correctness, and here the challenge is to explain away the impression that our judgements aim at absolute correctness. Mackie’s exotic answer is that even though there is no such absolute standard, the fiction that there is is nevertheless useful (Mackie 1977).

The general problem of non-objective judgement consists in a tension between (a) the inclination to regard these judgements as answerable to

objective standards of correctness, and (b) the difficulty to make epistemological sense of such standards. This problem is *especially* acute in the particular area of moral judgement. This is so because our inclination to regard moral judgements as answerable to objective standards of correctness is particularly strong, while the epistemological problems with moral facts are just as great as in other areas. In other words, the special problem of moral judgement stands out from the general problem in that a non-cognitivist answer is particularly problematic.

I just explained why non-cognitivism about moral judgements is harder to defend than non-cognitivism in other areas. The same is true of relativism: moral relativism is harder to defend than relativism in other areas, such as aesthetics. If moral relativism is the view that the standard of correctness for moral judgements can vary from thinker to thinker, then our everyday practice of arguing about moral matters seems to disconfirm moral relativism. Moreover, those advocating moral relativism may arouse the suspicion that they may be morally fickle or too tolerant towards morally outrageous acts.

In this paper, I want to outline an approach to the general problem and then explore a particular answer to the problem as it arises in the special case of moral judgements. The general approach involves the claim that the correctness of judgements quite generally should be regarded as relative to standards possessed by thinkers. However, as I will show, the approach nevertheless leaves room for large measures of objectivity in various areas. I shall be concerned with finding the right amount, and the right kind, of objectivity for moral judgements within the framework. Thus the ultimate aim of the paper is to sketch an account of judgement which (a) provides an answer to the question “Are moral judgements objective, and if so to what extent?” and (b) explains why moral judgements in particular present an especially difficult and controversial case. I shall start by summarising my approach to the general problem, then I shall develop and refine that approach and apply it to moral judgements.

#### THE GENERAL PROBLEM OF NON-OBJECTIVITY

I shall start by outlining my view about the general problem of objectivity by considering cases that are much easier and clearer than the case of moral judgement. Consider the following two pairs of sentences:

- (1) Snus is legal in Sweden.<sup>1</sup>  
 (2) Snus is rich in vitamins.  
  
 (3) Snus is enjoyable.  
 (4) *Göteborg's Rapé* is better than *General*.<sup>2</sup>

Sentences (1) – (4) share important features. They are all declarative. They express singular propositions and are similar in their grammatical form. They all express propositions, or contents, that it seems to make sense to call “true” or “false”. There are certain standards of correctness for each of them: one could err in believing the propositions expressed. Each has a negation, and the negation is incompatible with it. However, many people, including myself, differentiate between the first pair and the second pair in terms of the objectivity of their subject matter: (3) and (4) do not concern an objective matter while (1) and (2) do. (The reasons for saying that evaluative propositions like (3) and (4) lack objectivity are ultimately epistemological. For example: if we are to say that there are objective facts of the matter, to which our judgements on these matters are answerable, then we also need some account of how thinkers can access these very facts independently of one another.)

I follow Crispin Wright (1992) in thinking that objectivity or lack thereof is independent of truth-aptness. A proposition can be truth-apt yet lack objectivity. But what is objectivity if not aptitude for truth? According to the definition I shall be using, a proposition is objective just if we know in an a priori manner that if there is a disagreement on that proposition then at least one party of the disagreement is mistaken. Conversely, a proposition is not objective if a disagreement on it, for all we know in an a priori manner, may not involve any mistake. Let's consider the examples again. The idea is that a disagreement on (1) or (2) would show that one of the parties has made a mistake, while a disagreement on (3) or (4) would not. In other words, if you think Snus is rich in vitamins and I think it's not, then we know that one of us is wrong, and we know this in an a priori manner (more on this later). By contrast, if you think Snus is enjoyable and I think it's not, then that alone does not yet show that one of us is wrong.

<sup>1</sup> Snus is a type of tobacco popular in Sweden. It is used by placing it between lip and gums. Sweden is the only country in the European Union which permits (or fails to protect from) Snus.

<sup>2</sup> These are two well-known brands of Snus.

The matter is less clear for moral propositions, like the one expressed by:

(5) It is morally wrong to use Snus.

Here it is controversial whether (5) should be grouped with (1) and (2) or with (3) and (4) when it comes to objectivity. It is the ultimate aim of this paper to describe the kind of objectivity status enjoyed by judgements like the one expressed by (5) compared to that of those expressed by (1) – (4). In order to achieve this I will need to do a lot of stagesetting.

Mainstream theories of logic and language are quite hostile to the possibility of non-objectivity in the above-defined sense of a possibility of faultless disagreement. I have argued elsewhere (2002, 2003) that the best way of making room for this possibility is to relativise the truth of propositions (at least truth as it is relevant for semantics and logic) to something I called “perspectives”. This relativisation goes beyond the relativisation involved in possible world semantics, the main difference being that in the scheme I propose there is no analogue of the actual world, i.e. no uniquely privileged perspective. In the background of this relativistic semantics there is a wider conception of the relationship between thought, language and the world, which offers an explanation of why some propositions are objective others are not. In the next two sections I shall outline this wider conception.

#### CONCEPTUAL RULES AND A PRIORI KNOWLEDGE

The general conception in the background starts from the assumption that both thought and speech are social institutions. We frame thought in terms of *public* concepts, which are governed by social norms. It is crucial that there are these social norms, to which users of public concepts subject themselves, because this makes it possible for people to pool their cognitive resources. Because you and I both employ a common public concept of a traffic jam, and because there is a language we both know in which one can express the concept of a traffic jam, I can tell you that there is a traffic jam on the M25 *and* you can take advantage from receiving this information. There may well be exceptional occurrences of *idiosyncratic* thought, i.e. thought that is not framed in terms of public concepts, and perhaps conceptual progress

and innovation depend on it. However, in large part thought can be viewed as a social phenomenon.<sup>3</sup>

All the possible judgements within the repertoire of a competent thinker are subject to a number of norms. Any judgement is potentially correct or incorrect, because the activity of making judgements is constitutively governed by these norms. The process of becoming a competent thinker is the process of getting to know, and subjecting oneself to, these constitutive norms. In large part (but only in large part), this happens in childhood: an already competent thinker teaches a not-yet competent, or partially competent, learner. Teachers teach by example as well as by explicit instruction and explanation. Of course the only way in which teaching can take place is through the medium of a language: the processes of learning how to judge and learning how to use a language are intimately interdependent and largely coincide. For example, learning what the word “sleep” means will often coincide with learning the concept of sleep, its rules of application.<sup>4</sup> Knowing these constitutive norms of a concept will enable one to employ the concepts in making judgements. Often, making such judgements will proceed via a *posteriori* means, e.g. some evidential norm may oblige thinkers to apply a certain concept in the presence of a certain type of experiential input and thinkers may then use their sense-experience to arrive at judgements involving that concept. However, knowledge of constitutive rules can also give rise to a *priori* knowledge: a speaker will sometimes be able to arrive at a judgement without relying on experience merely by exercising non-perceptual abilities they acquired in the process of learning the concepts involved.

Some clarification and qualification is clearly needed at this point. First, being a competent thinker is not an all or nothing affair. From early childhood onwards we all gradually add more and more concepts or other tools of thought to our repertoires. The process is particularly

<sup>3</sup>This general approach is compatible with more individualistic conceptions of thought and language: it is possible to describe the same thinker in two different ways: as a thinker with his own unique system of thought and with his own unique idiolect, or as a thinker whose individual thought-processes participate in a communal system of thought and a communal language. However the approach that stresses the social aspects of thought will make it easier to describe certain social phenomena of thought and communication. In the current context this will be an advantage.

<sup>4</sup>Of course it is not impossible to master a concept without mastering any word that expresses that concept, but if the concept is a public one, then mastering it will typically involve learning to express it in a public language.

fast during a genetically programmed learning phase in early childhood. It makes sense to speak of a “mature thinker” when that phase has been completed. But it is neither easy to pinpoint the precise end of that process, nor is it easy (or possible) to spell out the exact conceptual range that an individual would need to master in order to count as a competent thinker. For example, while it seems obvious that the concept “sleep” is so basic that every competent thinker needs to master it, it is less clear whether every competent thinker grasps the concept of a frisbee, or of marmalade. But a competent thinker will be familiar with a critical mass of those concepts that are used in the community (or communities) of which he or she is a member. Secondly, the reliability and speed with which any given thinker can employ the concepts with which he or she is competent can vary. Thirdly, it may not always be clear mastery of which precise rules is necessary for competence with a given concept. It may or may not be necessary for possession of the concept of arthritis that one know that arthritis cannot occur in anyone’s thigh. (By contrast, knowledge that arthritis is a disease is presumably required.) I suspect that any concept can be mastered to varying degrees and that some thinkers, experts, have a socially privileged role in determining what is or is not regarded as constitutive of a concept. Because of the outstanding role of some experts, not every user of a given concept needs to be au fait with all the constitutive norms governing that concept. Some norms are not known by some of the users that are subject to them. Fourth, none of the above commits me to essentialism: it may be illuminating to view concepts as being governed by constitutive norms, but that does not entail that the objects answering to the concepts (if any) should be thought of as having some properties essentially. Nor am I committed to a conventionalist theory of necessity (see Boghossian 1996).

Rules constitutive of a concept may not always be expressible or explicable in words. Competence with a word can sometimes be achieved by exposure to examples and subsequent ability to apply the concept successfully. Thus, I may learn the concept of arthritis by experiencing what it feels like and a specialist telling me that what I am experiencing is arthritis. (Some people think that mastery of colour concepts even requires experiential familiarity with colours, i.e. those who say that blind people don’t have colour concepts. But I prefer to think of colour concepts as public concepts that blind people can use, even though they can’t use them in certain ways.)

The functioning of a community's system of thought and communication requires that there is a certain systematic convergence in the rules known or mastered by competent thinkers. The exact nature and extent of this convergence is too complex an issue to settle casually as an aside here. All that is important for my purposes here is that (a) concept use is governed by public norms, (b) sometimes experts may have a privileged role in determining which norms apply, and in those cases not all users, not even all competent users, may be aware of all the norms constitutive of a given concept. It is safest, therefore, to think of competence in thought as a matter of degree, where the degrees are measured on several scales: one can have a greater or smaller range of concepts in one's repertoire, one can have greater or smaller acquaintance with the norms governing a given concept and one can be more or less reliable and quicker or slower at applying the norms with which one is acquainted.<sup>5</sup>

I said earlier that knowledge of conceptual norms can give rise to a priori knowledge. I am now in a position to elaborate a little on that claim. Let's look at this from an individual's perspective first. One could say that an individual knows or believes a particular proposition in an a priori manner iff he or she has, in arriving at the belief, only exercised his or her general conceptual ability (i.e. the result of acquiring the concepts in question) without recourse to experience. On that individual notion of the a priori, some propositions are only accidentally known in an a priori way. For example, if my acquisition of the concept "red" involves experience of paradigm instances of red, my knowledge that these particular instances (which for me are paradigmatic) are red is a priori. It will not be a priori for other users who have acquired the concept in a different way. A more *communal* notion of the a priori would eliminate variation from individual to individual. One might characterize it as follows: a proposition is known/believed in an a priori manner just if it follows from the constitutive norms governing the concepts involved, so that any competent user of these concepts can know (correctly believe) it on the basis of conceptual norms only.

<sup>5</sup> In what follows, when I speak of competent thinkers, this is to be understood either as a vague reference to a mature thinker (i.e. one who has successfully completed the accelerated learning phase in childhood), or as relative to some particular ability, the context indicating which one. I.e. when I say that a competent thinker teaches an incompetent one some rules concerning the word "river", then I am speaking, obviously, of comparative competence with respect to those particular rules.

I have been talking about competence with sub-judgemental units: concepts. It is useful to think of judgements as being conceptually complex, and to view the ability to make any complete judgement as a complex ability which can be divided into various sub-judgemental, conceptual abilities. My ability to judge that the paint is dry can be usefully viewed as resulting from my ability to apply the concept of dryness, that of paint, and perhaps that of the binary quantifier expressed by “the”. A competent thinker will master a number of rules concerning each concept, and will also be subject to a number of norms. This sort of complexity explains why we can apply our conceptual abilities in new areas. It is also useful in the current context because it allows me to discuss differences in the kinds of rules concerning different kinds of concepts. For in the next subsection, I will attempt a classification of various concepts with respect to the kinds of rules governing them.

#### DIFFERENT WAYS OF CALIBRATING CONCEPTUAL ABILITIES

Each properly developed thinker is, as a matter of biological necessity, part of a community of thinkers and language users (at least at some time). For he or she must learn how to think and how to speak from other thinkers. But the dependence of individual thinkers on their community goes beyond our biological need to learn thought and language in a certain way. An individual's thinking faculties are designed to interact with a community of thinkers with whom the individual communicates. Each individual relies to a large extent on know-how and information acquired from others. An individual alone would not have the resources to sustain more than a very rudimentary system of thought. In my view, understanding these social aspects of thought is the key to understanding the special role that moral judgements play in our mental lives, and to resolving the special tension mentioned at the beginning of this paper.

It is hard to exaggerate the dependence of individual thinkers on the communities of thinkers of which they are members. Each individual will originally learn how to think and use language from others. Each individual acquires most of his or her knowledge from others. Each individual needs others to confirm that his or her own on-board faculties are functioning properly, i.e. that his or her senses are not deceiving him or her, that his or her memory is reliable, that his or her reasoning is correct. Each individual depends on the material help of others and



needs to cooperate with at least some. Each individual of course also feels a need for intellectual exchange with others.

A repertoire of common concepts as well as a common language are essential for this social interaction. How could we benefit from the testimony of others if we didn't employ the same concepts they do? How could we hope to influence others in their actions if we didn't partake in the same system of conceptual tools? It requires that members of the same community *calibrate* their conceptual abilities. There must be a range of shared concepts, concepts that each thinker employs following the same, or relevantly similar, methods. How am I to learn from your remark that there is a traffic jam on the M6 unless we share a concept of a traffic jam, i.e. use the concept according to the same rules? The idea of "sharing a concept" is ultimately a matter of being part of a community that enforces certain norms concerning that concept. For example, you and I may "share" a public concept of a traffic jam in that we largely agree on what can be correctly called a traffic jam, we both employ the same, or similar methods for establishing whether something is a traffic jam etc. But even where we do not agree we share the same concept because we are subject to the same norms and one of us may, according to these norms, have made a mistake.<sup>6</sup>

Successful calibration requires that different thinkers largely employ the same methods and follow the same rules because they are subject to the same rules. Being subject to a rule means being wrong when one violates the rule. Being wrong here is essentially a social idea. One is subject to the rule by being part of a community that enforces the rule, where enforcement involves social sanctions (see Brandom 1994, chs 1–3). It may be useful to distinguish two mechanisms of calibration: the teaching of novices and mutual adjustment among experienced users. In the process of the first, novices are introduced to conceptual rules by

<sup>6</sup> It is tempting to say that this need for calibration explains why humans have evolved so that they largely learn how to speak and think *from other members of their community*, rather than being born to develop their own idiosyncratic conceptual systems. There may be exceptions. Some humans are more or less idiosyncratic in their thought, some abilities are not learnt but innate. Universal grammar, according to Chomsky, constrains what languages can be acquired naturally. Perhaps there are also certain universal conceptual abilities with which we are born (for example Plato may have been right that we are born with the concept of equality or that of similarity. More relevant to the current investigation is the question whether he was right about our innate knowledge of concepts like that of justice!). However, the small extent to which conceptual and linguistic abilities may be innate or idiosyncratic cannot obscure the fact that most of our conceptual abilities are socially calibrated.

other, already experienced users. Since this is the standard route by which novices are introduced to new concepts, a degree of calibration is ensured right at the point of entry of new thinkers to a community of thinkers. But mutual adjustment among already experienced users also plays an important role in maintaining the community calibrated. Thinkers monitor one another and point out discrepancies in the norms followed. In conversation, they will often encounter divergences in opinion and following up such divergences will sometimes reveal a discrepancy in the conceptual norms the thinkers believe to be in force. In the case of teaching a novice a new concept, the teacher clearly enjoys a position of authority: the learner has to defer to the teacher's greater competence (except if the learner detects conflicts between the norms taught by the teacher and those norms with which the learner is already familiar). But in cases of mutual calibration of two already experienced users, one of them may or may not be in a position of authority—a mere “end-user”, for example, will defer to the authority of a frontline specialist.

The conceptual rules and methods that thinkers use and to which thinkers subject themselves, can take a variety of different forms. There may be rules that tie the applicability of a concept to certain experiential input. Thus a typical competent user of the concept “red” will reliably recognise red things on the basis of a certain kind of experience. Learning the concept will proceed by exposure to paradigm examples and foils until the learner reliably classifies the same things as red as his or her already competent teachers. Some may ask whether the experiences of learner and teacher are *really* the same and whether therefore they *really* follow the same rule. But there clearly is a kind of perceptual stimulus caused by red things that typical users respond to, and respond to very reliably. Whether that stimulus feels the same to all of them is beside the point. Successful calibration of a novice, in this case, requires that the novice achieve coincidence in judgement with his or her teacher(s).

Other conceptual rules may link concepts: “every bachelor is an unmarried man and vice versa” is perhaps such a rule. Teaching such a rule can take the form of an explicit definition. But it need not. One could teach the concept piecemeal, by explaining on separate occasions that being married is sufficient for not being a bachelor, that being male is necessary for being one, and that being male and never married is sufficient. It may also be possible to learn the concept of a bachelor just by examples and foils. However, it is doubtful that a user who cannot,

at least upon reflection, recognise that unmarried maleness is sufficient and necessary for bachelorhood has mastered the concept. (Suppose someone can reliably pick out bachelors from non-bachelors because he has been taught the piecemeal way and has picked up on a feature that all bachelors happen to share—let's suppose a certain laddishness, or a certain unkempt appearance etc. I believe they would be less than fully competent with the concept and ought to be set straight. This lack of competence might be revealed through discussion with others.)

Unlike “bachelor”, many other concepts do not admit of easy definition yet are not purely perceptual either. Consider “river”. Or the concept of flowing. Nevertheless teaching of such concepts succeeds: communities of thinkers are very well calibrated in their use of these concepts. Except in rare tricky cases, users reliably converge in their judgements on what counts as a river, for example.

In all the cases I have so far considered, the test for successful calibration is a convergence in the judgements made by different thinkers. It counts as the learner's failure if he or she applies the concept *red* (or *bachelor* or *river*) where the teacher would not. Success of calibration in these cases can be measured directly by the extent to which different thinkers coincide in their judgements, at least when no relevant uncertainty is involved. For want of a better word, let's call calibration of this sort “unisono” calibration. Not all concepts are calibrated in this way. Calibration of the concepts “is delicious” or “is prettier than”, for example, is not unisono but *polyphonic*. I fail to teach my niece the concept of deliciousness if she gets the impression that mastery of the concept of deliciousness involves judging the same things as delicious that I judge to be delicious. What my niece needs to pick up on is not just a certain quality of the food, objectively accessible both to her and to me, but rather she must learn to link judgements of deliciousness to her own practical reasoning. Deliciousness is an internalist concept: judging something to be delicious gives one a (pro tanto) reason to eat it. An individual who, in the absence of independent reasons against, prefers not to eat what he or she judges to be delicious, is irrational, i.e. fails to follow the norms of thought. Suppose you claim that the risotto is tastier than the hotpot. Then you choose the hotpot and not the risotto. Questioning reveals that you have no independent reason to choose hotpot over risotto which would outweigh the thought that risotto is tastier (e.g.: you believe the risotto will be hard to digest, or is more expensive, or you want to castigate yourself by depriving yourself

of culinary pleasures etc). This would show that you are either irrational or lack complete grasp of the concept of tastiness.

Corresponding to unisono versus polyphonic calibration is the distinction between objective and non-objective concepts. Concepts that are calibrated unisono are usually objective concepts: if two thinkers disagree on an objective matter, we know that one of them has made a mistake. By contrast, concepts that are polyphonically calibrated leave open the possibility of faultless disagreement. If my niece thinks *mortadella* is delicious and I disagree (i.e. believe it not to be delicious) then this does not (by itself) show that one of us is wrong. It makes sense to have non-objective concepts in addition to objective ones because it makes sense to have some internalist concepts with the kind of practical role that “delicious” has. It makes sense to have some concepts that are governed by constitutive rules which make the conditions under which it is correct to apply the concept depend on features of the thinker who is applying it.

The objectivity status of a concept, in my view, is largely an a priori matter. It is a matter for the social institution of thought and language to determine whether a concept’s role is that of an objective or a non-objective one. The status as objective or not of a concept is important, because this is what determines how a community deals with disagreements, i.e. how the thinkers of the community are calibrated. However, we can’t just stipulate objective properties into existence by deciding to treat certain concepts as objective. The world needs to play ball. Suppose a constitutive rule about the concept of deliciousness says that one can only correctly judge a thing to be delicious if it causes a certain kind of pleasure in one, or that judging it to be delicious gives us a reason to eat it. Suppose further that it turns out that the power of food to cause pleasure varies considerably between people. Then stipulating deliciousness to be objective would lead to problems. We would have to treat all disagreements about deliciousness as cases of error on the part of one of the disputants, but given the above-mentioned rule that links deliciousness to pleasure caused in people, we may not be able to find any mistake. Thus in my view, even though the status of a proposition as objective is an a priori matter, a system of thought and communication cannot go against the way the world is—or can do so only at the cost of severe problems.<sup>7, 8</sup>

<sup>7</sup> I have so far defined objectivity only for propositions, not for concepts or other sub-judgemental elements of propositions (Thanks to Wlodek Rabinowicz for pointing this

As in the case I just considered, there can be conflicts between the objectivity status and the methodological norms governing a concept. Consider again colour concepts. I believe we treat colour concepts as objective in the sense I have been using. Thus if I believe that this chair is red and you believe it is not, then at least one of us is making a mistake. However, there are also constitutive methodological norms concerning redness. We are all calibrated to classify as red or not red according to certain visual stimuli. Now if colour perception was not as stable as it is then this could not work. In that case, we would have to consider either adjusting the methodological norms or the objectivity status concerning colour concepts. (I suspect we would change the objectivity status.) But there are other areas of thought, where the objectivity status is central and will not usually be adjusted. Experimental science is such an area: scientific experiments must be repeatable, i.e. anyone following the same procedure must come to the same conclusion. When this is not possible, the conclusion is that the experiments and methods in question were unscientific. Thus we'll revise the methodological norms not the status of the relevant judgements as objective.

Making logical sense of the view I have sketched is not a trivial problem. Standard logical or semantical accounts do not leave room for non-objective thought, because they do not leave room for disagreements without error. I have argued elsewhere that the problem can be solved by relativising the notion of truth employed in semantics and logic. The same proposition can be true relative to one "perspective" and not true relative to another. Thinkers "possess" perspectives, and to judge or believe correctly is to judge or believe what is true according to one's own perspective. Perspective possession has properties that mirror our a priori rules – e.g. the perspective possessed by any thinker A and any thinker B will converge on objective propositions but might not converge on non-objective ones. Thus the apparatus of truth relativised to perspectives possessed by thinkers is just a formal model of the

out). Here is one proposal: a concept *C* is objective just if singular propositions containing it (e.g. of the form "*a* is *C*") are objective. This definition assumes that the source of all non-objectivity lies in the conceptual components corresponding to predicates, i.e. concepts, not in those corresponding to singular terms.

<sup>8</sup> It is meant to be part of this picture that conceptual rules can change over time (though not completely and all at once). Conceptual rules are not static but dynamic. Their gradual development is the outcome of complex social processes and negotiations.

system of a priori rules governing our interactions which I have just sketched.

DIAGNOSIS OF THE SPECIAL STATUS OF, AND  
TROUBLE WITH, MORAL JUDGEMENT

I have sketched an account of thought and communication that treats thought and talk as a social institution, and which makes room for a distinction between objective and non-objective areas of thought. What is the place of moral thought in this account? Are moral propositions objective or not, i.e. is faultless disagreement on them possible or not?

On the one hand, we again have an a priori normative link between the concept and whoever employs it (—at least according to internalists). Judging something to be morally wrong is ipso facto a subjective reason (pro tanto) not to do it. But different people may be guided by different systems of practical reasoning. What subjective reasons a person has or does not have is in part a matter of his or her personal constitution, preferences and history. Thus there is reason to say that faultless disagreement on moral matters is possible, so that moral propositions come out as non-objective.

On the other hand, there are a number of reasons why we are very reluctant to say that a disagreement on a moral matter could be faultless. Moral thought has an important co-ordinating function (see Mackie 1977, ch. 1.10, Wong 2005), and this seems to require that moral judgements be treated as objective and moral disagreements not be regarded as faultfree and legitimate. For this reason, moral matters cannot be discretionary in the way aesthetic matters are. We are happy to allow faultfree disagreements on whether this or that lampshade looks better, but much less happy to allow it on the question of whether torture is ever permissible. The reason for this is the co-ordinating function of moral thought. Moral judgements are practical and therefore guide our actions. But unlike other areas of practical deliberation, moral thought specifically concerns practical decisions that affect, or are affected by, others. If moral judgements were of the discretionary variety, then moral judgements could not coordinate. Thinkers would differ in their judgement of what ought to be done and there would be no reason to search for, and correct, any mistake. There would be no pressure ever to accept any inconvenient moral views, but such pressure (at least on some people) seems to be required for social co-ordination of action.

There is, therefore, a political reason for giving moral judgements objective status. Nevertheless, we cannot just conclude, for political reasons, that moral propositions are objective because after all, we still need to make epistemological sense of moral thought. The world needs to play ball. The rule, if any, that moral disagreements indicate error needs to harmonise with our rules for justifying moral norms, our methods of making moral judgements. However, there is a tremendous variety of different such methods across different cultures and groups. Moreover, in some cultures, and on some issues, there is notorious controversy surrounding the principles by which moral issues should be decided.

The key to understanding the situation is the fact that the political reasons for treating moral concepts as objective only have force within a certain sphere, i.e. within a certain *moral* community. The community of those with whom we aim to coordinate our actions for mutual benefit is relatively restricted. The community of thinkers with whom we can communicate extends much further than this restricted moral community. It is the fact that the moral community and the wider community of thinkers do not coincide which creates the mismatch between, on the one hand, our tendency to treat moral questions as objective but on the other hand failing to see how they could be objective because we can't make epistemological sense of it.

Let's consider, for one moment, a range of concepts that seem to be constitutively objective: the concepts involved in the description of scientific experiments. It seems to be part of basic scientific practice that disagreements at least about experimental data always count as an indication of error. If two scientists conduct the same experiment but they come to incompatible conclusions then some mistake must have been made in carrying out the experiment. In more familiar terms: scientific experiments must be repeatable. The repeatability of experiments seems to have a fundamental role in that experiments which are not are regarded as unscientific. In other words, the concepts employed in describing experiments and their outcomes only qualify as scientific if their evidential norms do not conflict with their status as objective. The result of this insistence that a given concept is objective (=that disagreements about it are always a sign of a mistake) is that the evidential norms concerning the concepts will develop in such a way that there is no tension with objectivity. For example, an objective concept of temperature emerged only once there were evidential methods that stand up to the principle that disagreements about temperature

show a mistake on the part of one of the disagreeing thinkers. In this case, I suppose the invention of a reliable thermometer was the decisive step. Thus, persistently treating a concept as objective will ultimately result in uniform and reliable evidential methods for that concept, which are themselves objective and guarantee that all who use the method correctly will arrive at the same conclusion.

Something similar will happen if a concept is treated as objective not for every thinker but merely within a community, i.e. if disagreements of members of that community count as showing there is a mistake. Thus within a moral community, the treatment of moral judgements as objective would lead to a certain methodological uniformity amongst users of moral concepts. But no such methodological uniformity will emerge beyond the boundaries of that moral community. This explains why there is, on the one hand, a temptation to regard moral judgements as non-objective—because there just seems so much methodological variety across a wider community of thinkers and no clear sense that much of the methodology is misguided. It also explains why, on the other hand, we are reluctant to regard moral judgements as non-objective or discretionary: to do so is to disregard the function moral judgements have within the communities in which they coordinate action.

In the current picture, we have a general community of thinkers and communicators, and this community may fall into smaller moral communities of people who share a set of principles or methods by which to decide moral matters. The special status of moral judgements arises because the required objectivity holds not across the entire community of thinkers but only across smaller moral communities. In fact, I believe, the picture is even more complicated. Some moral concepts command more general agreement than others. For example, it may be necessary to coordinate certain types of moral thought among friends and family but not among wider groups. It may be necessary to coordinate certain other areas of moral thought nationwide, but not internationally. Yet other moral norms may count as objective even in the widest possible community. The universal human rights seem to be an attempt to codify some such basic set of moral norms. It is possible, and probable, that individual thinkers can be members of many different social groups. It may be necessary to coordinate on moral values with some communities of which one is a member, but not with others.



REFINEMENT OF THE ACCOUNT OF OBJECTIVITY  
AND APPLICATION TO MORAL JUDGEMENT

We can now refine the notion of objectivity introduced earlier. A proposition may be maximally objective in the sense that any two thinkers' perspectives converge on it—or, put in terms of the possibility of faultless disagreement: a maximally objective proposition will be one on which there cannot be faultless disagreement among any pair of thinkers. A maximally non-objective proposition will be one on which any two thinkers' perspectives may diverge—or, put in terms of faultless disagreement: a maximally non-objective proposition is one on which faultless disagreement is possible among any pair of thinkers. Plausible examples are the proposition that there are 17 people in this room and the proposition that chocolate icecream is better than vanilla. Between these two extremes there are many shades of less than maximal and more than minimal objectivity. There is virtually unlimited room for differentiation. Objectivity among members of the same family, for example is the property a proposition has just if disagreement on it among members of the same family cannot be faultless. Objectivity among bee-keepers: the property a proposition has just if beekeepers cannot disagree on it without fault. For each symmetric relationship  $R$ , which can hold between two thinkers, we can define a refined notion of objectivity as follows: a proposition is  $R$ -objective just if disagreement among  $R$ -related individuals cannot be faultless.

The refined account permits a differentiated assessment of the objectivity of moral judgements. In the remainder of this section, I want to put forward some hypotheses as to what this assessment might be, but for reasons of space I will have to confine myself to the mere indication of a broad direction. I shall first try to assess whether moral judgements are maximally objective and whether they are maximally non-objective. Then I shall venture to identify some relations  $R$  such that some moral judgements turn out  $R$ -objective.

Are moral propositions maximally non-objective? Clearly not. It seems obvious that moral concepts are taught as objective among at least the mini-group of learner and teacher. Calibration is unisono. (Presumably, they are also taught as objective among a wider group, but I'll come to that later.)

Are moral propositions maximally objective? In order to keep the discussion more concrete, let's focus on one moral concept: wrongness. As I have said, wrongness is indeed objective among some groups (there is an  $R$  such that wrongness is  $R$ -objective). But would an a priori rule

of *maximal* objectivity harmonise with other constitutive norms concerning the concept of wrongness? Wrongness, at least for us internalists, is also a practical concept: if you think it's wrong then you have a reason not to do it, and in the absence of independent reasons in favour, doing it anyway amounts to irrational or incontinent action. I can see no reason why the internalist norm should conflict with maximal objectivity (thus on the current account, the standard argument from Humeanism and internalism to non-cognitivism, even if it works, has no immediate consequences for objectivity. Even non-cognitive states might be subject to objectivity).

Then there will be rules concerning the appropriate behaviour towards wrongdoers and the appropriate behaviour of the wrongdoers themselves. Thus, there may be a rule that it is morally permitted to sanction wrongdoers, or that it is wrong for wrongdoers not to regret/feel guilty/ashamed etc. These norms will vary between communities.

What about evidential norms? What counts as evidence for wrongness according to the a priori conceptual rules governing wrongness? Here opinions will diverge widely among philosophers. Moral theorists debate precisely what kinds of principles we should follow in making moral judgements, or even whether we should follow any principles at all. But let's free ourselves for a moment from the hang-ups of moral theory and assume that there is at least no straight route from the a priori norms governing the concept of wrongness to any particular ethical theory. How do ordinary people learn to use the concept of wrongness? I suspect they learn by a mixture of concrete cases and rules of thumb of limited generality. "Not burning your dead relatives: that's wrong." or "hurting others for no good reason is wrong." etc. This sort of teaching will lead to fairly good unisono calibration at least within the more immediate community of which the learner is a member. But evidential rules will usually vary across a wider community of which the individual is a part, and they will clearly vary considerably between disjoint communities.

Let us suppose that it makes sense to say that different communities employ the same concept of wrongness even though they follow different evidential rules.<sup>9</sup> Thus even though burying your dead parents may

<sup>9</sup> Two points: 1. what makes a concept a concept of wrongness? Clearly, if we are to say that morally quite diverse communities all share one concept of wrongness there must be some conceptual norms that are required for being a concept of wrongness. These might

not be evidence for wrongness in some Greek community while it is evidence for wrongness among some Indian group, both nevertheless are employing the concept of wrongness. Perhaps the internalist norm combined with the appropriateness of some social sanction (however varied) is sufficient for making a concept one of wrongness.

Does this tell us anything about whether wrongness is maximally objective? I believe it does. Consider the disagreement between the Greek who says burying your dead relatives is not wrong (in fact: required) and the Indian who says that burying your dead relatives is wrong. If we suppose wrongness is maximally objective, then it follows that one of them has made a mistake. One of them believes something that is not true on her own perspective. One of them would be better off, everything else being equal, changing her mind. This, I believe is an absurd conclusion. For all we know, both the Greek and the Indian may well have followed her own evidential norms correctly. Which norm is there that they have violated? In what sense could we say that they are indeed subject to that norm? Concluding that one must be wrong may also lead to explanatory problems. Assuming that the disagreement is not faultless would force anyone interpreting the behaviour of the two to diagnose a fault somewhere. But nothing in the cognitive and linguistic practices of the Greek's community may permit the conclusion that she is at fault, and nothing in the cognitive and linguistic practices of the Indian may permit the conclusion that she is at fault.

Therefore, if we are prepared to concede that two different communities with varying evidential rules can nevertheless both be employing the concept of wrongness, then we should also conclude that wrongness is not maximally objective.

Is it possible to draw any more specific conclusions about the kind of group amongst which moral concepts count as objective? There is reason to believe that for some social groups and relationships it is

be: the internalist norm (if it's wrong you have a reason not to do it) and perhaps some rules about guilt and shame. 2. Why don't we just say that different communities do not have the same concepts of wrongness, thus avoiding the whole issue? It does seem clear that there are, in many and diverse societies, concepts that roughly have the social co-ordinating function that moral concepts have in ours. Not drawing the connection would make our approach to other communities explanatorily barren—this is essentially the sociological argument for “methodological” moral relativism. However, if “employing the same concept” is understood to mean “being subject to exactly the same norms concerning that concept”, then members of these diverse moral communities clearly do not employ the same concept of wrongness.

constitutively important to share values. My conjecture is that, for example, the relationship of friendship requires agreement on a wide variety of moral issues: something, we assume, is wrong if friends disagree on a moral question (or perhaps a moral question of a certain sort). There are certainly some moral values discrepancies on which I can tolerate with non-friends, but not with friends. E.g. I would find it impossible to have a friend who is a racist. If we had a disagreement on a moral proposition in this area (“People of race XYZ don’t deserve the same rights as others”) then I would automatically take the disagreement to indicate that a mistake has been made and probably argue about the matter until we either agree or are no longer friends. There are other types of moral judgement on which I would probably be more tolerant of disagreement with friends—e.g. “Sex before marriage is wrong.” However, there may still be a narrower moral community within which disagreement even about these less important issues cannot count as faultless. Thus, within a family this may be so.

Thus my hypothesis is that moral concepts are objective, but not maximally objective, but rather objective within a group. Many are objective amongst friends, others are merely objective amongst family. More differentiation is probably needed.

#### BEYOND THE PALE PROBLEMS

A common criticism of moral relativists is that they end up unable to criticise those who commit wrongs but are not members of their own community.<sup>10</sup> In this final section, I want to argue, very briefly, that this charge is groundless and is probably the product of widespread acceptance of a certain universalism.

The charge is that according to the sort of moderate relativism I have sketched, I cannot criticise or act against people who are not part of my moral community, and who are committing acts that count as wrong according to the evidential norms in my moral community but not according to those in their own community. For example, I might want to act against members of a slave-owning society, who believe that it is morally permissible to keep slaves. The problem is that they may well

<sup>10</sup> The title of this section is an allusion to Harman’s doctrine (Harman 1975) that we cannot make moral judgements about those with whom we do not share a moral code, and that such people are therefore “beyond the pale”. For a more detailed discussion of this kind of criticism of moral relativism see Wong 1995 and 2005.

be making the right judgement according to their own norms, i.e. not making any sort of mistake. However, there really is no reason to believe that I cannot criticise them for, or act against, their practices. I can act against them because my own moral judgements require me to do so. I can criticize their actions because they are wrong by my own lights. Nothing in the account prevents that.

The following may be the argument on which the criticism is based. Moral norms are norms of rationality. So, anyone is subject to moral norms which, were they fully rational, they would be able to see. In particular, if it was morally right for us to prevent them from carrying on with their slave-owning practice, then they, being fully rational, would have to concede that this is so. But relativists think that it is possible that our slaveowners are not committing any mistakes, that they are making fully competent and correct use of the moral concepts of their community, but are nevertheless not recognising our moral obligation to intervene. If we were right, they, being rational, would see it. But they don't. So we are not right.

I am inclined to deny the first premise, namely that moral norms are universal norms of rationality. I do not, at the moment, see how the argument would work without that premiss.

It may be true that, in some sense, I cannot criticise them for their views: the scenario dictates that they are exercising their thinking faculty fully competently. However, I may well believe that they are employing a system of moral concepts which is in itself worthy of criticism, and I could criticise them at this meta-level. Their system may be inconsistent or it may conflict with some undeniable moral claim. If, however, I do not have any such basis for criticism, then there might well be a situation in which I cannot offer any rational criticism. At this point, all I can do is persuade them by non-rational means (by luring them over to my way of viewing things, or perhaps by brute force). I believe that in this sense, relativism may well entail that "one cannot criticise" the slaveowners of the example. However, in this sense I believe it to be an acceptable consequence.

REFERENCES

- Boghossian, Paul 1996: "Analyticity Reconsidered". *Nous* 30, pp. 360–91.
- Brandom, Robert 1994: *Making it Explicit*, Cambridge, Mass.: Harvard University Press.
- Harman, Gilbert 1975: "Moral Relativism Defended". *Philosophical Review* 84, pp. 3–22.
- Kölbel, Max 2002: *Truth without Objectivity*, London: Routledge.
- Kölbel, Max 2003: "Faultless Disagreement". *Proceedings of the Aristotelian Society* 104, pp. 53–73.
- Kölbel, Max 2004: "Indexical Relativism vs Genuine Relativism". *International Journal of Philosophical Studies* 12, pp. 297–313.
- Mackie, John L. 1977: *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- Wright, Crispin 1992: *Truth and Objectivity*. Cambridge, Mass.: Harvard University Press.
- Wong, David 1995: "Pluralistic Relativism". *Midwest Studies in Philosophy* 20, pp. 378–99.
- Wong, David 2005: "Taking Moral Relativism Seriously". This volume.