

Moral realism, disagreement, and conceptual ethics

Michael Klenk 

Institute for Values and Philosophy, Delft University of Technology, Delft, The Netherlands

ABSTRACT


Moral disagreement is often thought to be of great metaethical significance for moral realists. I explore what remains of that significance when we look at moral disagreement through the lens of a combination of two influential and independently plausible hypotheses about moral language. The Morality-As-Cooperation (MAC) hypothesis says that our capacity for and use of moral language is an adaptation to increase mutualistic cooperation. The Concepts-As-Tools (CAT) hypothesis says that we often engage in disputes about language use and that many apparent moral disagreements are linguistic disagreements in disguise. The combined MAC-CAT view that I explore suggests that we frequently engage in linguistic disputes to find optimal means for mutualistic cooperation. I show that this perspective weakens sceptical claims based on moral disagreements, that is offers a novel way for moral realists to explain the apparent genuineness of moral disagreements without the need to accept theses borrowed from non-cognitivism.

ARTICLE HISTORY Received 7 June 2021; Accepted 25 August 2021

KEYWORDS Moral disagreement; conceptual ethics; moral anthropology; moral scepticism

1. Introduction

On the conceptual ethics picture, concepts are like tools, and using them has downstream effects.¹ For example, considering something RAPE can have moral, psychological, legal, social, and economic effects for the involved parties. Given what is at stake in concept use, proponents of the conceptual ethics picture suggest that we often engage in debates about language. They also argue that many apparent disagreements about evaluative and normative matters may turn out to be *primarily* disagreements about language. So, looking through the lens of conceptual

CONTACT Michael Klenk  m.b.o.t.klenk@tu-delft.nl  Institute for Values and Philosophy, Delft University of Technology, Jaffalaan 5, Delft 2628BX, The Netherlands

¹Given a context; sometimes the stronger claim is made that they have specific functions.

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

ethics, moral disagreements disappear and re-emerge as linguistic disagreements. This claim bears metaethical interest because the phenomenon of moral disagreement has long been thought to raise difficult problems for moral realism David, 1986.²

In this paper, I explore the implications of the conceptual ethics view on the metaethical significance of deep moral disagreement. My aim is to show that the conceptual ethics view afford moral realists a novel reply to the semantic disagreement challenge (which I will explain below). The plan is as follows. The next three sections introduce the topic of this paper, providing the necessary details on the disagreement-based challenge for realism, as well as the Morality-As-Cooperation (MAC) and Concepts-As-Tools (CAT) hypotheses. I offer critical clarifications, but I will be brief in these sections as they describe arguments defended elsewhere. My contributions to the debate come in the sections that follow. I present the combined MAC-CAT view and explain how it defuses the disagreement challenge (sections 4 and 5). Then I present evidence for the MAC-CAT view from recent (cultural) anthropology and conclude (6 and 7).

The paper's innovation lies in exploring the connection of two influential and independently plausible hypotheses about moral language for the evaluation of deep moral disagreement as a central metaethical phenomenon. Its findings should be of interest both for metaethicists and proponents of either hypothesis.

I am primarily focused on defending the conditional that *if* the MAC and CAT theses are true, then realists have a novel answer to the semantic disagreement challenge. For any apparent deep moral disagreement that initially seems damning for the realist's view (as suggested by the semantic disagreement challenge, explored in the next section), this paper implies that we must consider that we have a semantic, meta-linguistic disagreement about how best to facilitate cooperation instead, which does not have sceptical implications. This is the novel and additional line of defence against the semantic disagreement challenge for robust realists.

The scope of my claim is thus that there are at least *some* cases of apparent deep moral disagreement that can be explained in that

²Moral realism is the theory that (at least some) moral truths are constitutively and causally independent of human attitudes or beliefs and that (at least some) moral knowledge is possible (Brink 1989; Railton 1989; Shafer-Landau 2003; Enoch 2011). 'Objectivity' is usually cashed out in terms of the mind-independence of moral facts. This conception of realism is encompassing, by including both non-naturalism and naturalist versions of realism. It does not, however, presuppose a continuity between morality and other fields of inquiry, like science or mathematics.

fashion.³ For the record, I believe that MAC and CAT imply that *most* if not *all* cases of disagreement can be explained in that way. But that is not my claim in this paper. All that is needed to secure its aims is to show how MAC and CAT imply a novel escape route for realists in at least some cases. The paper thus contributes to a deeper theoretical understanding of moral realism and its relation to the theses about morality-as-cooperation, and conceptual ethics.

2. A disagreement-based semantic challenge for moral realism

This section introduces descriptive relativism, deep moral disagreement, and the semantic challenge against moral realism that is based on it.

Descriptive ethical relativism is the view that, as a matter of empirical fact, people's moral convictions differ drastically across different geographies, societies, cultures, or periods. Undeniably, some peoples engage in practices that other cultures, such as contemporary Europeans or Americans, would find objectionable. For example, consider some marriage practices and genital mutilation today, and cannibalism and human sacrifices in Roman or Aztec cultures (cf. Prinz 2007). Hence, not only students in introductory philosophy classes but many prominent anthropologists (Westermarck 1932; Benedict 1934; Evans-Pritchard 1937) and philosophers believe it to be evident to anyone with an elementary understanding of the history and cultures of the world that descriptive ethical relativism must be true (though see Klenk [2019] for a assessment of the recent cultural anthropological evidence that weakens the claim of descriptive relativism).

A *deep moral disagreement* is a moral disagreement that the disputing parties cannot rationally resolve insofar as they exclude from consideration the grounds that lead them to endorse the disputed claim.⁴ All parties endorse on proper grounds what their subjective evidence supports and yet they disagree. Anti-realists about morality have sometimes argued that many moral disagreements are deep and that this poses a problem for moral realism on epistemic or metaphysical grounds.

³The scope is restricted further to actual rather than possible moral disagreement because, as will become clear below, the semantic challenge involves an inference to the best explanation. That would seem, epistemically, more forceful explaining actual or at least 'nearby' phenomena (to wit, those that happen in nearby possible worlds). The greater we imagine the moral disagreement to be, the less it would appear to be located in a nearby possible world Klenk (2018a). Hence the focus on actual moral disagreements.

⁴The requirement to bracket the disputed claim in assessments of a disagreement is known as the 'independence' requirement in the literature on peer-disagreement.

The anti-realist argument that I focus on in this paper invokes moral disagreement against moral realism through an explanatory claim (Stevenson 1937; Gibbard 1990; Hare 1991; Tersman 2006; Horn 2019). According to the argument, moral realists cannot explain the sense of genuine disagreement we feel when we encounter a moral disagreement. Consider Hare’s famous cases of a tribe of cannibals that uses the term ‘good’ as a general term of praise, but unlike a missionary visiting the tribe, they apply the term to very different people and actions than the missionary. While the missionary applies the term ‘good’ to, say, gentle and compassionate people, the cannibals use it to praise rash and aggressive people. Nonetheless, argues Hare, the missionary and the cannibals still mean the same when using the term ‘good’. Because if they would not, argues Hare (1991):

... then when the missionary said that people who collected no scalps were good (English), and the cannibals said that people who collected a lot of scalps were good (cannibal), they would not be disagreeing, because in English (at any rate missionary English), “good” would mean among other things “doing no murder”, whereas in the cannibals’ language “good” would mean something quite different, among other things “productive of maximum scalps”.

Hare considers this conclusion patently absurd. To save the appearance of moral disagreement, Hare proposes a non-cognitivist analysis of moral language, which allows us to locate the moral disagreement in the conflicting plans or preferences of the disputing parties. The most elaborate recent version of Hare’s argument is by Tersman (2006). I loosely reconstruct his version of the semantic challenge (what he calls the ‘argument from ambiguity’) as follows:

Semantic challenge (first pass)	
MORAL COGNITIVISM	Moral cognitivism is true.
MORAL DISAGREEMENT	There is moral disagreement (and so disputing parties follow different rules when applying the disputed term).
CHARITY	If moral cognitivism is true and disagreeing parties follow different rules when applying the disputed term, they mean different things by that term.
VERBAL DISPUTE	If they mean different things, then they cannot be in disagreement.
CONTRADICTION	So, there is moral disagreement (assumption), and there is no moral disagreement (from MORAL COGNITIVISM, CHARITY, and VERBAL DISPUTE). ⁵

We have to reject one of the premises to avoid the contradiction in the conclusion. Assuming that the CHARITY premise is true (cf. Tersman 2006)

⁵If there is no disagreement, then, by extension, there is no moral disagreement.

and given the existence of moral disagreement, the contradiction seems best resolved by rejecting moral cognitivism because then the appearance of moral disagreement can be retained.⁶

The argument is not complete yet, because moral realists could attempt to escape the contradiction by rejecting the VERBAL DISPUTE premise. Moral realists could accept that the disputing parties *mean different things* and yet that they are in moral disagreement because of a clash of attitudes, plans, or preferences. But when realists explain moral disagreement in virtue of the conflicting plans or preferences expressed, they rely on considerations that are external to their view (Tersman 2006). Moral realism would be parasitic on the explanatory power of non-cognitivism, and the better explanation seems to be that moral non-cognitivism is true. So:

Semantic challenge (continued)	
EXPLANATORY OPTIONS	Disputing parties mean different things and still disagree because of a clash of attitudes (moral cognitivism) OR they mean the same and therefore disagree (moral non-cognitivism).
BEST EXPLANATION	Moral non-cognitivism is explanatorily superior.
CONCLUSION	So, the explanatorily superior route to avoid contradiction is moral non-cognitivism.

So, in a nutshell, realism (via cognitivism) initially gains plausibility points because it can explain why moral disagreement is pressing (explanation: it concerns the truth of a proposition), but if moral disagreement rides on differing meaning of a term, then a charitable interpretation is that the disputing parties in those disagreements do not literally express incompatible content and so the realist’s explanation is undermined: moral disagreements do not concern the truth of a proposition.⁷ Moral non-cognitivism comes out as explanatorily superior compared to cognitivism. Given its commitment to moral cognitivism, the argument

⁶Of course, one could challenge the CHARITY premise by developing novel semantic moves by realists. Tersman (2006) provides an extended discussion and rebuttal of these moves (see also Fisher 2007). Enoch (2009) also discusses the semantic challenge (and sets it aside *as a challenge from disagreement*), but not the CHARITY premise in detail. Notwithstanding the possibility of challenging the CHARITY premise, I will accept it for the sake of argument and explore a realist answer to the argument *even if* the premise is accepted.

⁷The semantic challenge is not supposed to deal a death-blow to moral realism Tersman (2006). It is not the case that moral realism fails because moral cognitivism fails, with the phenomenon of moral disagreement making the difference as the explanandum. Nothing as strong can be concluded from the semantic challenge. Nonetheless, moral disagreement is supposed to provide plausibility to moral cognitivism. These plausibility points are now gone, the argument goes, and moral non-cognitivism gains some. But cognitivism may still stand for other reasons and thus realism may stand as well. My paper here extends the discussion and a line of thought begun by Enoch (2009) who raises some issues along these lines.

goes, moral realism is weakened.⁸ That is the disagreement-based semantic challenge to moral realism.

However, I will now argue that moral realism need not incur an explanatory loss *at all*. Actual moral disagreements may be practical disagreements about how to best cooperate in disguise. This possibility comes into view when we take two recent popular hypotheses about the function of morality and the nature of disagreements seriously. I will briefly introduce both theses, and then show what follows for the semantic disagreement-based argument against moral realism. My introduction of both hypotheses will mostly cover familiar territory, and readers may want to go straight ahead to section 4, where I show how both hypotheses complement each other.

3. The Morality-As-Cooperation hypothesis

According to the Morality-As-Cooperation (MAC) hypothesis, our capacity for and use of moral language – representation of norms and values – is an evolutionary adaptation to foster mutualistic cooperation. Barkhausen already puts the morality-as-cooperation hypothesis in explicitly linguistic terms (Barkhausen 2016):

A major function of moral language and the cognitive capacities required for its use is to coordinate on mutually beneficial rules of conduct in situations with scope for common gain.

There is much recent philosophical support for the MAC hypothesis (Kitcher 2011; Sterelny and Fraser 2016), and several philosophers invoke it in metaethical arguments (Joyce 2006; Sauer 2019).

The most impressive support for the MAC hypothesis to date comes from a series of studies conducted by Oliver Curry and colleagues. Curry (2016) looks at strategies that yield mutualistic gain in several non-zero sum formal games and hypothesise, in line with MAC, that behaviour that is in accordance with those strategies (e.g. reciprocal altruism, conflict resolution by signalling) will be evaluated as morally good (and behaviour that isn't will be evaluated as morally bad). Based on initial support for their hypothesis (Curry, Chesters, and van Lissa 2019b), Curry, Mullins, and Whitehouse (2019a) they tested the

⁸The semantic challenge assumes that moral realism needs moral cognitivism. Moral realists may be able to reject that assumption Kahane (2013). I think that this is a valid strategy, insofar as realists can make sense of knowledge without belief Radford (1966); Myers-Schulz and Schwitzgebel (2013). However, most moral realists in fact defend accounts that would be hurt by undermining moral cognitivism, and so I set this strategy aside here.

hypothesis against extensive ethnographic records on 60 societies from around the globe. They looked for indications of moral evaluation (e.g. use of the terms good, bad, virtuous, moral, a duty, etc.) and assessed whether the described behaviour or character trait fell into one of six categories of cooperative behaviour (helping family members, helping group members, engaging in reciprocal cooperation, being brave, respecting superiors, sharing a disputed resource) (Curry, Mullins, and Whitehouse 2019a). They assessed 962 cases, and in 99.9%, cooperative behaviour had positive moral valence (*ibid.*).

In light of the arguments advanced by other proponents of MAC, and Curry's recent quantitative work, there is good reason to think that moral language, as a representation of norms and values, is an evolutionary adaptation to increase mutualistic cooperation and I will assume it in what follows. Moreover, it is also important to note, for later purposes, that people are predisposed for cooperative behaviour as shown in several studies from developmental psychology (e.g. Tomasello 2016). There is good reason to think that humans are 'cooperation machines'. We should, therefore, expect that people, in general, are somewhat disposed toward seeking mutualistic cooperation.

Nonetheless, we started out acknowledging the intuitive pull of descriptive ethical relativism. There is much anecdotal evidence that people are engaged in moral disputes and thus that moral disagreement is real. The next section shows what we can say about this from the perspective of another hypothesis about moral language.

Before proceeding, it is worthwhile to ask whether the MAC hypothesis is supportive of moral realism in the first place.⁹ Evolutionary explanations of morality in general have in recent years been considered a problem for moral realism.¹⁰ If MAC indeed helps realists to evade the semantic disagreement challenge, but undermines moral realism in other ways, then a realist defence against the semantic challenge may be but a pyrrhic victory. But whether that is the case is far from clear. First, even if successful, the evolutionary challenges levelled against moral realism may not be knock-down arguments of the theory. It will thus still be relevant whether realists can evade a challenge on another front. Second, and more to the point, it is doubtful that the familiar evolutionary

⁹Thanks to an anonymous referee for suggesting this objection.

¹⁰For instance, Joyce (2006) suggested that evolutionary explanations undercut the claim that there are moral properties (realistically conceived), and Sharon Street (2006) at some point suggests that evolutionary influence on our moral beliefs would make them, at best, coincidentally true in problematic ways. These are but two ways to interpret recent anti-realist arguments based on evolutionary explanations of morality. See Klenk (2017, 2018b); for further discussion.

challenges against realism succeed. Amongst other things, it is debated just why, and if at all, evolutionary explanations of morality would undercut the justification of moral beliefs (Klenk 2022).¹¹ At this stage, the negative relation of evolution and moral realism remains inconclusive. With these considerations in mind, it is interesting to explore how evolutionary considerations may ultimately turn to support rather than undermine realism. In order to do that, I will now turn to the second hypothesis explored in this paper.

4. The Concepts-As-Tools hypothesis

Moral disputes serve as indicators that there are moral disagreements, and we identify moral disputes when we observe people following different rules in applying the same moral term.¹² People apply terms like ‘torture’ to different things, or show conflicting behaviour toward things labelled with the same term, like ‘abortion’. It seems prudent to follow the proponents of the semantic challenge and conclude that these are cases of moral disagreement.

However, the Concepts-As-Tools (CAT) hypothesis says that we sometimes discuss language, as a symptom of underlying other disputes, and that our disputes about language are genuine disagreements nonetheless. CAT gains support from work on conceptual ethics. Plunkett and Sundell have, in several papers, provided compelling reason to think that moral disagreements may primarily be disagreements about language and not about morality (Plunkett and Sundell 2013, 2019; Plunkett 2015). On their view, we often (aware or not) engage in discussions about how to use a term (what they dub ‘conceptual ethics’) because of the downstream effects of using a term. They distinguish descriptive linguistic disagreements (about how a term is in fact used) and normative linguistic disagreements (about how a term should be used). The latter is relevant in what follows because it explains why moral disagreements seem like genuine disagreements. Alas, what is at stake is not an object-level issue about morality, but a linguistic issue about how a term ought to be used. So, some disagreements may be genuine, and speakers ‘feel opposed’ toward each other, even if there is no disagreement in the literal content of what they express.

¹¹Evolutionary considerations may, from a virtue epistemic perspective, even strengthen the moral justification of our moral beliefs, cf. Klenk (forthcoming).

¹²Nothing should hinge on the term ‘rule’ here – I just mean that people are sufficiently consistently using terms differently.

Metalinguistic negotiation matters because linguistic choices have practical consequences, as the introductory remarks about apply the concept RAPE illustrated. In particular, metalinguistic negotiations matter when there are shared expectations about what a term implies in practical terms. They suggest that we should expect normative linguistic disagreement in two cases: in discussions of semantic standards at play in a given context that govern the application of gradable adjectives (e.g. tall) and in discussions about the concepts that should be associated with a term (e.g. athlete). Specifically:¹³

Given a certain social-historical setting—a setting in which certain words (largely independent of which specific concept they express) fill specific and important functional roles in our practices – participants might care a great deal (and genuinely substantive results could hang on) which concept/word pairings we employ in a given context. (Plunkett and Sundell 2013)

Davies ([forthcoming](#)) is clear about the situations in which we should expect metalinguistic negotiations: situations where the truth of a sentence has consequences. When we want to avoid torture, it matters whether we classify a situation as torture. We can disagree about the labelling because we agree about the possible outcomes of assenting or denying the labelling. That is why Plunkett and Sundell (2019) write that metalinguistic negotiation would not make much sense when we disagree about what to do about torture.

The passage suggests that we care about concepts as tools: using them (in the sense of associating them with a term) has real-world consequences that the speaker might care about deeply. But that does not mean that they are doing so consciously – speakers do not need to be aware when they consider concepts as tools for things they care about. So, if we accept CAT, we should always check whether some apparent moral disagreement is not, in fact, primarily a linguistic disagreement.

5. MAC-CAT: disputing language for cooperative benefit

Combine the MAC and CAT views and you get the following picture:

MAC-CAT: In order to achieve cooperative benefits (from MAC), we engage in linguistic disagreements (from CAT).

The MAC-CAT view is attractive because each hypothesis complements the other. The MAC view leaves us with a question about the processes

¹³See also Davies ([forthcoming](#)).

by which people coordinate mutualistic behaviour. Turning to CAT, we learn that concept- and term use can fulfil that function (in some cases).¹⁴ The CAT view leaves us with a question about what metalinguistic negotiation is a symptom of – why do we engage in such disputes? Turning to MAC, we learn that we engage in metalinguistic negotiation because we are in disagreement about how to best realise mutually beneficial cooperation.

On MAC-CAT, our linguistic disagreements are practical disagreements about how to best achieve cooperative benefits. The scope of the MAC-CAT view is broad. For all apparent cases of moral disagreement, we need to check whether there is a moral or a linguistic, practical one. We should distinguish three broad classes of normative linguistic disagreements here that Plunkett and Sundell do not sufficiently keep apart:¹⁵

Second-order normative linguistic disagreement: We agree descriptively about the consequences of using a term T, but we disagree about whether we should have these consequences. In that case, our disagreement about T is a placeholder of our deeper normative disagreement about the desirability of the consequences. For example, two people might agree that applying the term ‘torture’ to a practice will make that practice banned by law, but they might disagree about whether we should ban torture or not.

Second-order descriptive linguistic disagreement: We agree normatively about the consequences of using T (e.g. that all practices labelled ‘torture’ should be banned) but we disagree whether these consequences will, in fact, occur. In that case, our disagreement about T is a placeholder for our deeper descriptive disagreement about consequences.

Futile disagreement: We agree neither normatively nor descriptively about the consequences of using T.

According to CAT, we should not expect futile disagreement and so, for each linguistic disagreement we encounter, we only need to determine whether it is a second-order descriptive- or normative disagreement. Do the disputing parties normatively agree about the consequences? If so, they will be in a second-order descriptive disagreement: they disagree about the use of a term because they are in disagreement about whether using the term will achieve a common goal. According to MAC, we should expect that people have a common goal: reaping cooperation benefits.

¹⁴Note that the evidence here is rather thin and anecdotal – as in the introductory example about the concept RAPE. I assume that these hunches can be corroborated by further studies.

¹⁵One could add a further distinction about people’s evaluation of term T as a means to reach consequences C. People might agree both descriptively (that T will lead to C) and normatively (that C is good), but disagree that T should be used for C.

So, MAC-CAT suggests that we should expect lots of second-order descriptive disagreements, rather than second-order normative disagreements. People, according to MAC-CAT, negotiate language-use as a tool for cooperation.

MAC-CAT seems immediately suspect because it presupposes that people will share common ground about the aims of interaction and, more specifically, that they aim at mutualistic cooperation. That might seem especially worrying since we are considering moral disagreement, quite the opposite of common ground. That seems to exclude situations where there is less common ground about, for example, the expected or correct consequences of doing something.

There is considerable scope to push back against this worry. First, people seem deeply sensitive to the effects of language-use, as indicated by political movements aimed at ‘problematism’ (Keane 2016). More importantly, the MAC view helps us to the claim that people have an extraordinary proclivity toward pro-social, cooperative behaviour (cf. Tomasello 2016). It is thus not too far-fetched to think that they share a common aim of enhancing cooperative behaviour. Unless we have reason to suspect that people are *not* in a situation in which both realise that there is potential for mutualistic benefit, we should suspect their disagreements, according to MAC-CAT, to represent practical disputes about how to make cooperation go best.

6. The MAC-CAT view against the semantic challenge

This section aims to argue that MAC-CAT offers a novel response to the semantic challenge. Assuming MAC-CAT, we should expect that what appears to be a moral disagreement really is a disagreement about tool-choice disagreement in disguise.

MAC-CAT offers a new perspective on apparent moral disagreements. It implies that disputants may *in the first instance* do not disagree about morality on an object-level, but about how to use language and we should often expect them to share a common goal and be in agreement about the desired consequences of using language (from MAC). That opens up a path for moral realists to avoid the contradiction and the explanatory loss generated by the semantic challenge. They can reject VERBAL DISPUTE, which falsely implies that a lack of sameness in meaning is sufficient to rule out disagreement. Even if people mean different things by a moral term (following CHARITY and moral cognitivism), they can still be in a disagreement, and the explanation of this does not

favour either cognitivism or non-cognitivism. Realists can thus accept a charity principle about translation according to which consistently different rules for application indicate that a term is used with different meanings, and yet explain why moral disagreements are disagreements without being parasitic on the explanatory power of non-cognitivism.¹⁶

MAC-CAT thus takes the pressure out of the disagreement-based semantic challenge against moral realism because, if the view is accurate, it is no longer clear that the *clash of attitudes* explanation is the best explanation of the sense of genuine disagreement that many experience when they think about moral disagreements. Instead of the clash of attitudes, the disagreeing parties might also just experience a *practical* disagreement about how to best achieve mutualistic cooperation, and they are discussing which ‘tools’ to use to achieve that end. Anyone who has engaged in practical disagreements – how to cook a steak at a garden party, whether to use crampons or not on a snowy hillside, how to set-up an organisation for success – knows that they can be fierce, even though the disputing parties are aiming for a mutual goal. So, the fact that disputing parties feel opposed can be explained by their practical disagreement, and we need not invoke the non-cognitivist idea that they are not expressing beliefs at all (and attitudes instead).

However, the MAC-CAT perspective on the semantic challenge comes under pressure from the thought that metalinguistic negotiation is a surface phenomenon that indicates some deeper disagreement that is moral after all. As Plunkett and Sundell themselves admit, linguistic disagreements is *a symptom* of some moral disagreement down the line. So, when two people disagree about the concept associated with the term ‘athlete’, they are having a metalinguistic disagreement (a practical dispute, on the MAC-CAT view), but only *primarily* and *not purely* a practical dispute, because somewhere down the line they probably have a substantive moral disagreement, for example about whether animals can be the subject of praise and blame (as an athlete presumably is). With that admission on the table, the pendulum of explanatory power may just swing back toward non-cognitivism.

Deriving negative consequences from that admission for realism is premature and speculative. There might be a moral disagreement down the line, but it just might – MAC-CAT does not require that this is the case. We

¹⁶One might also argue that it is a requirement that the moral disagreement persists to discharge the explanatory burden of the semantic challenge. This seems suspect to me because it is not clear that moral disagreement persists on the non-cognitivist account, either. In any case, discussing this is beyond the scope of this paper.

just do not know whether the disagreement sticks around when the linguistic disagreement is settled.

Of course, there will always be a concrete question about classifying a given, particular moral disagreement as a moral or a linguistic one. For each apparent moral disagreement, we need to start out assuming, following MAC-CAT, that we face a linguistic disagreement, and we should expect that it is a second-order descriptive one. In that case, however, we have replaced a (potentially deep) moral disagreement with a practical one and thereby avoid contradiction (per the semantic challenge) without requiring resources beyond moral realism. Exploring this possibility by combining the two theses – MAC and CAT – is the central innovation of this paper. Moreover, I have shown that the MA-CAT view is a coherent combination of two plausible hypotheses, and stronger than both theses inspected individually because either thesis answers open questions that plague the other.

My point is not that this dissolves the semantic challenge. I have not argued that many moral disagreements are linguistic ones, let alone provided a substantial defence of MAC or CAT. Nonetheless, I believe I have made plausible the hitherto unacknowledged claim that the combination of MAC and CAT, two independently plausibly hypotheses, opens up a new possibility for realists to offer ‘defusing’ explanations of moral disagreements (or if you want to play that game, require anti-realists to show that an apparent moral disagreement indeed is not a linguistic disagreement).

7. Corroborating the argument and extending its scope

Above I have argued that the MAC-CAT view is a coherent combination of two plausible hypotheses about moral language. If it is true, we should expect at least some apparent moral disagreements to be linguistic disagreements. While the above introduces a novel reply to the semantic disagreement challenge for realists, I will close by briefly considering how the MAC-CAT thesis can gain further independent support.

An important way to corroborate the combined MAC-CAT view would be to evaluate whether actual cases of alleged moral disagreement are well explained as linguistic disagreement. That is, whether the disagreement is primarily practical (as predicted by MAC-CAT) or whether it bottoms out in moral disagreement (thus blocking the MAC-CAT-based reply to the semantic challenge).

Some might think that this is hopeless because they are convinced by descriptive ethical relativism that there is (a lot of) genuine moral disagreement. Its existence has long been, and in some circles still is, accepted as a fact of anthropology. Take Mackie (Mackie 1977 emphasis added):

The argument from relativity has as its premise the well-known variation in moral codes from one society to another and from one period to another, and also the differences in moral beliefs between different groups and classes within a complex community. Such variation is in itself merely a truth of descriptive morality, *a fact of anthropology* which entails neither first-order nor second-order ethical views.

The ethical turn in cultural anthropology should shake up the alleged fact of anthropology (cf. Klenk 2019). Michelle Moody-Adams, in a perceptive review of anthropological research before the ethical turn, claims that ‘it is difficult (at best) to establish that one has indeed found a genuine instance of fundamental disagreement’ (Moody-Adams 1997). Laidlaw remarks that the sources used to establish the fact of anthropology are a ‘rather eccentric election’ (Laidlaw 2014). Thus, though descriptive ethical relativism has been taken as an established fact for long, recent developments in philosophy and empirical fields have tended to yield defusing explanations of alleged moral disagreement. The historical fact that philosophers have overestimated the extend and depth of moral disagreement, and now recognise several ‘defusing explanations’ of moral disagreement should make us open to the possibility that MAC-CAT will further continue the diffusion of allegedly worrisome moral disagreements. In recognition of this possibility, the MAC-CAT view should be corroborated by detailed qualitative assessments of alleged moral disagreements. Recent work in cultural anthropology (e.g. Keane 2016) allows us to inspect a wider range of richly described cases, which can be one testing ground for the MAC-CAT view.

In studying these cases, we should expect, from the MAC-CAT perspective, that people fundamentally agree about moral values insofar as they were evolutionarily relevant, and that they disagree about which application of concepts – and corresponding behavioural effects – will best serve their shared values. Proponents of MAC-CAT need not claim that *all* apparent moral disagreement turns out to be linguistic disagreement. As suggested above, it is enough for them to show that some allegedly deep moral disagreement, which may give rise to the semantic challenge for realism, may turn out to be linguistic and thus unproblematic. Thus, it

should initially suffice to show that for some cases, the MAC-CAT prediction is borne out. One thus needs to ascertain whether people indeed hold shared values that correspond to mutualistically beneficial outcomes and whether the apparent moral conflict indeed can be interpreted as a linguistic negotiation about how to realise them. In analysing whether that is the case, one need not presuppose that people see their disagreements in this light. People engaged in seeming moral disagreements may very well interpret them as a genuine moral disagreement, and their claims and arguments as making moral rather than linguistic or practical points. That is compatible with MAC-CAT. The CAT view does not imply anything about reflective access to the practice of negotiating terms, and the MAC view suggests that being unreflective about the function of moral language may be required or at least conducive to its function to be fulfilled.

The individual support of MAC and CAT along with the trend toward defusing explanations of moral disagreement may incline us to expect the MAC-CAT view corroborated by such empirical studies. Until they have been produced, however, the MAC-CAT view remains a hypothesis, albeit one that represents an intriguing novel possibility realists to respond to the semantic disagreement challenge and thus to escape one of their most difficult tests.

8. Conclusion

In this paper, I have explored the combination of two interesting hypotheses about moral language, Morality-As-Cooperation (MAC) and Concept-As-Tools (CAT), and put it to work in a novel context: the semantic disagreement challenge against moral realism. I think that the implications of MAC-CAT are plausible if we consider that many of our moral disagreements seem to involve a common goal hindered by a confusion of tongues – we sometimes just fail to decipher each other's moral code. With MAC-CAT supported by the plausibility of each individual hypothesis, the additional strength gained by combining both, and the further support it may gain by explaining well some recent cases from cultural anthropology, I believe that the paper suggests a novel realist-friendly route for assessing the semantic challenge for realism. Rather than being in tune with recent evolutionary-based challenge against moral realism, this paper also indicates how evolutionary explanations of morality may turn out to support moral realism. Future work may fruitfully explore whether cases of apparent moral disagreement, documented

for example in the cultural anthropology literature, may be well explained by the MAC-CAT view, to further strengthen its support.

Acknowledgements

I received helpful comments on an earlier draft of this paper from participants at the 2020 MetaEssen IV workshop at the University of Duisburg-Essen. I should acknowledge, in particular, Peter Königs, Woo Ram Lee, Hichem Naar, Neil Roughly, Eleonora Severini, and Folke Tersman. I also received helpful comments on that version from Richard Rowland, and a later version benefitted from discussion at the 'Foundations & Synthesis' research colloquium of the Ethics of Socially Disruptive Technologies research programme in May 2021. Guido Löhr provided detailed and helpful written comments on that version, too. Finally, I am grateful to the referees for this journal for their comments.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

Michael Klenk's work on this publication was part of the project ValueChange that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 788321.

ORCID

Michael Klenk  <http://orcid.org/0000-0002-1483-0799>

References

- Barkhausen, Max. 2016. "Reductionist Moral Realism and the Contingency of Moral Evolution." *Ethics* 126: 662–689.
- Benedict, Ruth. 1934. "Anthropology and the Abnormal." *The Journal of General Psychology* 10 (1): 59–82. doi:[10.1080/00221309.1934.9917714](https://doi.org/10.1080/00221309.1934.9917714).
- Brink, David. 1986. *Moral Realism and the Foundations of Ethics*. Vol. 95, 163–207. Oxford University Press.
- Curry, Oliver Scott. 2016. "Morality as Cooperation: A Problem-Centred Approach." In *The Evolution of Morality*, vol. 43, edited by Todd K. Shackelford, and Randal D. Hansen, 27–51. Cham: Springer.
- Curry, Oliver Scott, Matthew Jones Chesters, and Caspar J. van Lissa. 2019b. "Mapping Morality with a Compass: Testing the Theory of 'Morality-as-Cooperation' with a

- New Questionnaire." *Journal of Research in Personality* 78: 106–124. doi:[10.1016/j.jrp.2018.10.008](https://doi.org/10.1016/j.jrp.2018.10.008) .
- Curry, Oliver Scott, Daniel Austin Mullins, and Harvey Whitehouse. 2019a. "Is It Good to Cooperate? Testing the Theory of Morality-as-Cooperation in 60 Societies." *Current Anthropology* 60 (1): 47–69. doi:[10.1086/701478](https://doi.org/10.1086/701478) .
- Davies, Alex. *forthcoming*. "Identity Display: Another Motive for Metalinguistic Disagreement." *Inquiry* 64 (8): 861–882.
- Enoch, David. 2009. "How Is Moral Disagreement a Problem for Realism?" *The Journal of Ethics* 13 (1): 15–50. doi:[10.1007/s10892-008-9041-z](https://doi.org/10.1007/s10892-008-9041-z) .
- Evans-Pritchard, Edward E. 1937. *Witchcraft, Oracles, and Magic among the Azande*. Oxford: Oxford University Press.
- Fisher, Andrew. 2007. "Review of 'Moral Disagreement' by Folke Tersman." *Notre Dame Philosophical Reviews*.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings. A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Hare, Richard Mervyn. 1991. *The Language of Morals*. Oxford: Clarendon.
- Horn, Justin. 2019. "Does Moral Disagreement Pose a Semantic Challenge to Moral Realism?" *Philosophia*. doi:[10.1007/s11406-019-00142-z](https://doi.org/10.1007/s11406-019-00142-z) .
- Joyce, Richard. 2006. *The Evolution of Morality. A Bradford Book*. Cambridge, MA: The MIT Press.
- Kahane, Guy. 2013. "Must Metaethical Realism Make a Semantic Claim?" *Journal of Moral Philosophy* 10 (2): 148–178. doi:[10.1163/174552412X628869](https://doi.org/10.1163/174552412X628869) .
- Keane, Webb. 2016. *Ethical Life. Its Natural and Social Histories*. Princeton, NJ: Princeton University Press.
- Kitcher, Philip. 2011. *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Klenk, Michael. 2017. "Can Moral Realists Deflect Defeat Due to Evolutionary Explanations of Morality?" *Pacific Philosophical Quarterly* 98: 227–248. doi:[10.1111/papq.12207](https://doi.org/10.1111/papq.12207) .
- Klenk, Michael. 2018a. "Evolution and Moral Disagreement." *Journal of Ethics and Social Philosophy* 14 (2). doi:[10.26556/jesp.v14i2.476](https://doi.org/10.26556/jesp.v14i2.476) .
- Klenk, Michael. 2018b. "Survival of Defeat: Evolution, Moral Objectivity, and Undercutting." PhD thesis, Utrecht University. <https://www.narcis.nl/publication/RecordID/oai:dspace.library.uu.nl:1874%2F364788>.
- Klenk, Michael. 2019. "Moral Philosophy and the 'Ethical Turn' in Anthropology." *Zeitschrift für Ethik und Moralphilosophie*, 1–23. doi:[10.1007/s42048-019-00040-9](https://doi.org/10.1007/s42048-019-00040-9) .
- Klenk, Michael. 2022. "How to Defeat Objectivist Moral Beliefs." *American Philosophical Quarterly* 59 (1): 43–60.
- Klenk, Michael. *forthcoming*. "Virtue Epistemology and Evolutionary Debunking." In *Evolutionary Debunking Arguments*, edited by Diego E. Machuca, 1–29. New York, NY: Routledge, available online at <https://ssrn.com/abstract=3897649>.
- Laidlaw, James. 2014. *The Subject of Virtue. An Anthropology of Ethics and Freedom*. New York, NY: Cambridge University Press.
- Mackie, John L. 1977. *Ethics. Inventing Right and Wrong*. London: Penguin.
- Moody-Adams, Michele M. 1997. *Fieldwork in Familiar Places. Morality, Culture, and Philosophy*. Cambridge, MA: Harvard University Press.

- Myers-Schulz, Blake, and Eric Schwitzgebel. 2013. "Knowing That P Without Believing That P." *Noûs* 47 (2): 371–384. doi:[10.1111/nous.12022](https://doi.org/10.1111/nous.12022) .
- Plunkett, David. 2015. "Which Concepts Should We Use?: Metalinguistic Negotiations and the Methodology of Philosophy." *Inquiry* 58 (7-8): 828–874. doi:[10.1080/0020174X.2015.1080184](https://doi.org/10.1080/0020174X.2015.1080184) .
- Plunkett, David, and Tim Sundell. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13 (23): 1–37.
- Plunkett, David, and Tim Sundell. 2019. "Metalinguistic Negotiation and Speaker Error." *Inquiry* 50 (2): 1–26. doi:[10.1080/0020174X.2019.1610055](https://doi.org/10.1080/0020174X.2019.1610055) .
- Prinz, Jesse J. 2007. *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- Radford, Colin. 1966. "Knowledge – By Examples." *Analysis* 27 (1): 1–11. doi:[10.2307/3326979](https://doi.org/10.2307/3326979) .
- Sauer, Hanno. 2019. "The Argument from Agreement: How Universal Values Undermine Moral Realism." *Ratio* 32 (4): 339–352. doi:[10.1111/rati.12233](https://doi.org/10.1111/rati.12233) .
- Sterelny, Kim, and Ben Fraser. 2016. "Evolution and Moral Realism." *The British Journal for the Philosophy of Science* 14: axv060. doi:[10.1093/bjps/axv060](https://doi.org/10.1093/bjps/axv060) .
- Stevenson, Charles Leslie. 1937. "The Emotive Meaning of Ethical Terms." *Mind; A Quarterly Review of Psychology and Philosophy* XLVI (181): 14–31. doi:[10.1093/mind/XLVI.181.14](https://doi.org/10.1093/mind/XLVI.181.14) .
- Tersman, Folke. 2006. *Moral Disagreement*. Cambridge: Cambridge University Press.
- Tomasello, Michael. 2016. *A Natural History of Human Morality*. Cambridge, MA: Harvard University Press.
- Westermarck, Edward. 1932. *Ethical Relativity*. Westport, CN: Greenwood Press.