

A DISTINCTION WITHOUT A DIFFERENCE? GOOD ADVICE FOR MORAL ERROR THEORISTS¹

Hallvard Lillehammer

Abstract

This paper explores the prospects of different forms of moral error theory. It is argued that only a suitably local error theory would make good sense of the fact that it is possible to give and receive genuinely good moral advice.

1. Error theories and skeptical puzzles

One aspect of our philosophical tradition consists in illuminating different areas of thought by investigating skeptical puzzles arising within these areas of thought. The underlying aim of this project is not best understood as either establishing or refuting the relevant skeptical conclusion, as opposed to (for example) articulating the various assumptions on which claims within these areas of thought depend. Some recent work on moral skepticism has taken a different form. Thus, some recent arguments for moral skepticism have been taken by their proponents to present a genuine challenge to the credentials of moral thought as such, in the absence of a response to which all moral thought has been said to rest on some kind of error or mistake. This paper is a response to that trend which, I argue, is based on an overestimation of the philosophical significance of certain manifestations of analytical semantics and ontology. The claim is not that there are no skeptical arguments capable of throwing some aspects of moral thought into doubt. Instead, the suggestion is that plausible skeptical arguments in metaethics will necessarily be local.

2. Moral Error theories

Let an error theory about a set of claims, S, be a theory asserting that the best way to interpret the claims in S is as aiming at truly or otherwise correctly representing some feature of the world, and that there is no adequate interpretation of any of the ('atomic', or otherwise basic) claims in S on which they come out as either true or correct in the relevant sense.² For example, it has famously been argued by J. L. Mackie and others that moral thought is best interpreted as being committed to the existence of mind independent, intrinsically motivating and reason giving entities; that there are no such entities in the world as described by natural science; and that moral thought is therefore committed to a metaphysically erroneous picture of the world.³

¹ I am grateful to the audience at the 2012 Ratio Conference at Reading University for helpful comments and criticisms of the talk on which this paper is based. I am also grateful to Niklas Möller, and to various participants in Ethics Group and Metaethics seminars at Cambridge, for insightful discussions of related issues.

² For the significance of the qualification, see Charles Pigden, 'Nihilism, Nietzsche, and the Doppelgänger Problem', in R. Joyce and S. Kirchin (eds.), *A World Without Values: Essays on John Mackie's Moral Error Theory*, (Dordrecht: Springer, 2010), pp.17-34.

³ See J. L. Mackie, *Ethics: Inventing Right and Wrong* (Harmondsworth: Penguin, 1977) and Richard Joyce, *The Myth of Morality* (Cambridge: Cambridge University Press, 2001).

There is a distinct kind of moral error theory definable for every set of moral claims that someone could be suitably skeptical about. Unsurprisingly, therefore, moral error theories have not always had universal scope.⁴ In particular, there is a set of interlocking strands in modern philosophy that have regarded either ‘absolute’ or ‘categorical’ (as opposed to ‘conditional’ or ‘hypothetical’) obligation; ‘external’ or ‘necessary’ (as opposed to ‘internal’ or ‘contingent’) reasons; or ‘thin’ or ‘deontic’ (as opposed to ‘thick’ or ‘evaluative’) concepts as good candidates for some kind of localised error-theoretic, skeptical, or otherwise deflationary, treatment.⁵ What these strands have in common is the acceptance of a vantage point within moral thought broadly understood (such as the domain of ‘the ethical’, as defined by Bernard Williams) from which it is asked whether we can make coherent sense of some further subset of moral claims in light of other things we believe about the world. On this kind of approach, the error theoretical challenge has traditionally been formulated against the background of the prior acceptance of other moral claims that are taken to stand on a comparatively firmer footing, but which are thought to fall short in attempts to capture ‘hard cases’ (such as Hume’s discussion of the ‘Sensible Knave’), or to reconcile particular aspects of our intellectual tradition with a secular form of philosophical naturalism (such as Nietzsche’s critique of Christian morality). In this paper I make no attempt to assess the plausibility of these kinds of localized moral error theory. Whatever their ultimate (de-) merits, I shall assume that at least some error theories of this kind are at least minimally plausible, and therefore that investigating the implications of accepting some kind of moral error theory is a coherent and empirically tractable project.⁶ My question is whether this is the only kind of error theory that has any prospect of being both plausible and interesting. My guiding hypothesis is that the answer is affirmative.

Philosophical arguments for and against different kinds of error theory are complicated by fact that at least some moral claims are capable of more than one coherent (even if partly revisionary) interpretation. In its standard formulation, an error theory about some set of claims, S, entails that there is no acceptable interpretation of any of the claims in S on which any of these claims come out as true or otherwise correct (or as not purely aiming to represent some aspect of the world at all). Given this formulation, it is not enough to establish an error theory that there is *some* interpretation of the claims in S on which they all come out as false, incorrect, or otherwise mistaken. As is well known, the appearance of a claim in the form of a declarative sentence (e.g. ‘That’s

⁴ Not even Mackie’s error theory is obviously best interpreted as being universal in scope (See Mackie, *Ethics*, pp. 25-7; 105ff). Examples of universal error theories include views according to which all normative claims, whether moral or non-moral, are either false, incoherent or otherwise mistaken in some way. See e.g. Bart Streumer, ‘Can We Believe the Error Theory?’, *The Journal of Philosophy*, forthcoming.

⁵ Other examples of localized error theory would be theories targeted at a subset of moral claims that presuppose some problematic conception of individual freedom, or the extent to which different moral claims command reasonable convergence. See e.g. David Wiggins, *Needs, Values, Truth* (Oxford Blackwell, 1982) and Bernard Williams, *Ethics and the Limits of Philosophy* (London: Fontana, 1985).

⁶ Thus, it might be argued that we ought to withdraw some or all of the problematic claims; explicitly or implicitly relativize them; treat them as ‘fictional’; or construe them as products of contingent social construction. For further discussion, see Hallvard Lillehammer, ‘Constructivism and the Error Theory’, in C. Miller (ed.), *The Continuum Companion to Ethics* (London: Continuum, 2011), pp. 55-76.

terrible!') could serve to express a variety of different attitudes in different contexts, from a reflectively sophisticated attempt to give a complete and cardinal ranking of a range of practical options on the one hand to a primitive expression of outrage or disgust on the other. Given the variety of minimally plausible interpretations that have actually been given of sincere expressions of personal attitude involving moral vocabulary, one might wonder if the problem here is not so much whether there are too few interpretations around to make at least some moral claims come out as true or otherwise correct, but rather that there are too many.⁷ In what follows, I mainly ignore the implications of this possibility - with one important exception. This is the possibility that what is currently being observed in metaethical debate is partly a substantially normative (indeed, in some cases a substantially moral) discussion about the 'ownership' of different sets of variously interpretable vocabularies. I shall return to this issue shortly.

So what, if anything, determines whether an interpretation of S (whether error theoretical or not) is good enough to capture what goes on when someone gives sincere expression to a moral claim? In this paper, I make one fundamental (and obviously not theoretically innocent) assumption about the necessary conditions for an adequate interpretation of the moral claims in S, as actually employed by competent speakers in a given context of application. I assume that an adequate interpretation of S is constrained by a plausible description of actual moral deliberation and agency, as this is experienced by individuals who sincerely and competently employ moral concepts. I take this to require that the interpretation in question succeed in taking the moral deliberation of the individuals involved seriously as distinctively moral (as opposed to, for example, narrowly prudential, purely manipulative, or otherwise insincere). In making this assumption, I do not assume that the interpretation in question must be first-personally transparent. Nor do I assume that the self-ascription of an interpretation on the part of a subject is 'factive', or even plausible. In other words, it would be possible for us to make accurate sense of what someone is doing when sincerely expressing a moral claim even if that person is not, in their actual situation, able to make accurate sense of themselves as doing what we take them to be doing (such as when, as a victim of false consciousness, I explicitly endorse an exclusionary practice that in fact oppresses me).

When I say that an interpretation of S is 'good enough', I mean to leave open the following possibility, aptly described by David Lewis in the course of his articulation of a dispositional theory of value:

I suggest that (for some of us, or some of us sometimes) the ... dispositional theory best captures what it would take for something to perfectly deserve the name 'value'. There are no perfect deservers of the name to be had. But there are plenty of imperfect deservers of the name, and my... version is meant to capture what it takes to be one of the best of them... What to make of the situation is mainly a matter of temperament... When it comes to deserving a name, there's better or worse but who's to say how good is good enough?⁸

⁷ For example, the availability of constructivist and response dependent accounts of moral truth generate the prospect of an indefinite number of interpretations of moral claims on which at least some of those claims are likely to be true, and (in some cases) necessarily so. See Lillehammer, 'Constructivism'.

⁸ David Lewis, 'Dispositional Theories of Value', *Proceedings of the Aristotelian Society*, Suppl. Vol. 63 (1989), pp. 139-174.

The discussion of moral error theories in this paper is consistent with the basic point that Lewis articulates in this passage. Where I depart from Lewis is in explicitly emphasizing that the choice between different interpretations can itself be a moral issue (and not, therefore, ‘a matter of temperament’, as that phrase is normally understood). In other words, my aim is to draw explicit attention to the fact that the question ‘How we should speak about the values?’ is one the answer to which ought to be informed by our best evaluative judgement. In this paper I also stop short of defending a version of the dispositional theory of value (moral or otherwise). Although the claims made in this paper are consistent with a dispositional theory of value, they do not entail it.

3. Moral uncertainty, epistemic asymmetries and good advice

Sometimes people are genuinely uncertain what to do. Sometimes there are less uncertain people around who are able to give these uncertain people good advice. In some such cases, the uncertain person stands in a relationship of epistemic asymmetry vis-a-vis their advisor that puts that advisor in a comparatively privileged position to say what the uncertain person ought, or ought not to do (or what it would be better or worse for the uncertain person to do). This can be very fortunate, both for the people in question and for the rest of us, especially if the uncertain person in question is seriously considering whether to do something wrong or bad. I regard this as a ‘datum’ that any philosophical theory ought to accommodate; not a questionable assumption we could reasonably deny.

I say that an epistemic asymmetry occurs with respect to some issue when subject X is better placed than subject Y (who could be X at some different time or place) with respect to the answer to some question Q regarding that issue (where Q could be either a ‘practical’ question of how to act or a ‘theoretical’ question of how to think about some issue relevant to Q). I say that an epistemic asymmetry is moral when X is better placed than Y with respect to Q, and Q is a moral question (e.g. about what it is right to do, or what it would be best to do, etc.). In the case of good moral advice there are at least two questions in play. The first is whether or not it is right (or best, etc.) for X to adopt some candidate course of action or other (such as speaking out in the course of a meeting). The second is whether or not it is right (or best, etc.) for X to take Y’s advice, whatever that might be (such as ‘Better keep your mouth shut’). My focus in what follows is primarily on the second question, although I shall also briefly comment on the first, and on the relation between the two.

Epistemic asymmetries abound in moral thought. They include cases of cultural learning (as when someone is in danger of trusting the wrong people, or failing to grasp the moral significance of offers or threats); personal vulnerability (as in crises of confidence, or being vulnerable to implicit biases or self-deception); moral divisions of labour (as in managing confidential information, or developing role specific social or personal insights); and crises of conscience (as where someone finds that their moral assessment of a situation is out of line with the assessment of a person or institutional representative that has a more comprehensive, or otherwise superior, view of the matter at hand). They also include bog standard cases of moral uncertainty, as in cases where one person is unable to make up their mind about the morally best course of action available, but is fortunate enough to receive the advice of another person with a clearer or more sensitive view of the issue (as in cases where someone takes moral advice from a trusted, and trustworthy, friend).

Some epistemic asymmetries in moral thought boil down to asymmetries in

evidence on issues that are not themselves moral. Thus, I may be uncertain about whether or not to speak out about some issue in the course of a meeting because I'm not sure what my future career prospects will be if I do. Other epistemic asymmetries are arguably moral all the way down, as in the case where my uncertainty about whether or not to speak out boils down to uncertainty about the relative values of integrity, loyalty and self-interest in the context of professional relationships.⁹

All cases of epistemic asymmetry in moral belief take place against a background of existing personal, social, institutional and historical expectations and commitments.¹⁰ Thus, if I am uncertain about whether to speak out during a course of a meeting, any satisfactory answer to that question should be sensitive to facts about me, such as my prior commitments, as well as facts about my situation, such as the norms and rationale of the institution from within which I ask the question, and the practices or traditions of which that institution is a part. In some cases, I might ask the question of whether to speak out bracketing all these facts, as though I were trying to decide the matter from 'nowhere', or from 'the point of view of the universe'. Yet in most cases, I will take at least some of these facts, including some of my prior commitments, for granted in deciding what to do. Furthermore, in some cases this will not only be responsible; it will obviously be right. Consider the following, minimally moralized, example by way of analogy: there might be no 'absolute', 'unconditional', or otherwise 'ultimate', foundation for someone's choice between different vegetables to buy in the supermarket on the way home from work. Indeed, in abstraction from a given set of contextual parameters the question could be largely indeterminate. This does not mean that there are no more or less responsible ways to choose between available options within those parameters; much less that any moral, prudential, or otherwise normative judgement made against the background of such parameters would be false, incorrect, or otherwise mistaken. Any metaethical theory that implies the opposite would give an implausible description of how practical thinking actually works. Even worse, it could result in unreasonable prescriptions for how it ought to work.

Epistemic asymmetries can obviously be more or less difficult to overcome. In some cases (such as institutional acculturation), it could simply be a matter of time or opportunity. In other cases (such as cases of deeply held prejudice or mental illness), the obstacles may be difficult, if not practically impossible, to surmount. There are also cases where X, who is epistemically better placed than Y with respect to some moral matter, is nevertheless some way short of comprehensively grasping the moral significance of the case at hand. Although this case raises interesting questions of its own (e.g. about how, if at all, X ought to advise Y in such a case), it does not follow that it is never sensible to rely on someone else's moral advice, when that advice is misleading, or strictly speaking incorrect. Not all misleading or strictly speaking incorrect beliefs lead away from the truth. And when they do not, it can be perfectly sensible to rely on them.

Talk of epistemic asymmetries between moral beliefs might raise eyebrows among those who would deny either that there are any moral beliefs; that any moral beliefs could be justified if *these* moral beliefs are false; or that any moral beliefs could be justified if *all* moral beliefs are false. To this kind of skepticism there is a simple line

⁹ The point is controversial, but nothing I say in this paper hinges on it.

¹⁰ For a more textured discussion of a case of moral uncertainty and epistemic asymmetry in the context of a crisis of conscience experienced by a senior government official, see A. I. Applbaum, *Ethics for Adversaries: the Morality of Roles in Public and Professional Life* (Princeton: Princeton University Press, 1999), pp. 209ff.

of response. By all means say that there are no moral beliefs *strictly speaking*, or that false moral beliefs cannot be justified *strictly speaking*. With respect to the question at hand, we are entitled to the use of a vocabulary in terms of which to describe the fact that it is possible for someone to find themselves in a state of moral uncertainty where there are more or less sensible ways to respond, and where some people are in a better position to say what these more or less sensible ways of responding are. If we are forced by our love of theoretical simplicity to use other terms than ‘epistemic asymmetry’ or ‘moral belief’, then we can always make some new terms up. As I will argue in the following section, however, we should resist this kind of conciliatory response in the case of epistemic asymmetries in moral belief.

Be that as it may, the main point is that in some circumstances it is obviously right, best, or just sensible to take moral advice from someone who is comparatively better placed with respect to a moral issue about which one is, oneself, uncertain. In such cases, refusing to take account of the advice of the person who is comparatively better placed could be a wrong, worse, or a silly thing to do. To deny that it is ever sensible to be guided by good moral advice, or that there really is such a thing, is deeply implausible. It is deeply implausible because it fails to make interpretive sense of the predicament of those who actually find themselves at either end of such advisory relationships; who seriously wonder what to make of them; and whose decision on whether to give or take advice can be both a moral and an epistemic achievement.

The cases of good moral advice described in this section are a sub-set of cases involving choice under normative uncertainty.¹¹ Some of the recent literature in moral philosophy has tended to draw our attention away from this phenomenon, partly by focusing on thought experiments where all elements of risk and uncertainty are deliberately factored out (as in standard versions of the Trolley Problem), or by focusing on cases of moral disagreement where participants enter the discussion with their moral beliefs already formed.¹² Yet many insightful and morally important claims are made in conditions of considerable uncertainty. Any plausible account of moral thought has to account for such insights as a central feature of social life. Some metaethical theories (including universal moral error theories) would classify all moral claims made in such cases as literally false or incorrect. In so doing, they would classify as mistaken or erroneous a set of accurate descriptions of such cases as involving a grasp of genuine moral insights. To that extent, such metaethical theories should seem as unattractive as they are implausible.

4. Putting error theories in their place

Given its reliance on substantively moral descriptions of social reality, the discussion in the previous section is structurally similar to a number of well-known arguments against skeptical hypotheses in ethics that start from the assumption of some 'Moorean' (and

¹¹ Clearly, there could be epistemic asymmetry without moral uncertainty, as well as moral uncertainty without epistemic asymmetry. I have nothing further to say about these possibilities here.

¹² For further discussion of the philosophical significance of disagreement in morality and elsewhere, see e.g. R. Feldman and T. A. Warfield (eds.), *Disagreement* (Oxford: Oxford University Press, 2010).

therefore unquestionable) moral fact.¹³ Where the discussion in the previous section departs from at least some ‘Moorean’ arguments is in its comparative neutrality with respect to the moral facts of the case about which the agents involved are ultimately concerned (e.g. whether or not I really ought to speak out at the meeting). Thus, the success conditions of the moral claims involved in such cases do not depend on the advisors in question actually knowing what those facts are. Instead, the relevant success conditions relate to the people at either end of an epistemic asymmetry primarily as responsible moral deliberators.¹⁴

This distinctive feature of the discussion in the previous section might be thought to have at least two major disadvantages. The first is that we arguably have no conception of what good moral advice is in a given situation apart from our conception of what it is, in fact, right (or best, etc.) to do in that situation. The second is that in focusing on the example of good moral advice, the discussion in the previous section fails to make use of any of the more robust, and obviously plausible, moral claims (about the wrongness of cruelty and so on) on which ‘Moorean’ arguments against moral error theories have traditionally focused. I basically agree with both criticisms, but offer the following observations. First, the discussion in the previous section is obviously consistent with the soundness of existing versions of the ‘Moorean’ strategy that take a more conventional form.¹⁵ Second, I believe the case of good moral advice is sufficiently distinctive and interesting in its own right to merit separate attention in this context. On the one hand, it draws attention to the fact that at least some moral deliberation is a socially embodied activity during the course of which various normative commitments and parameters are already in play. Second, it draws attention away from a focus on particular moral claims and the ‘Is there really a (single) right answer?’ or ‘How do you know that?’ questions they sometimes provoke. Much that is deep and insightful in moral thought takes place in circumstances where we don’t have an answer to either or both of those questions (or, were we are lacking in confidence that we have one). The example of good moral advice draws attention to this fact without getting bogged down in the details of any particular moral issue.

Another issue on which the discussion in the previous section departs from some ‘Moorean’ arguments is that it does not imply the reduction of the question of moral truth or correctness to a first order question about our entitlement to affirm a set of substantial moral claims. Thus, it does not follow from this discussion that there is no such thing as a coherent second order activity of semantics or ontology the success conditions of which are (to some extent) independent of first order moral thought and practice. What does follow from this discussion is that the only sense in which such an activity could show

¹³ The caveat ‘broadly’ is appropriate because nothing in the previous discussion entails that a ‘Moorean’ strategy needs to involve a commitment to the existence of ‘obvious’, ‘self-evident’, or ‘axiomatic’ moral truths.

¹⁴ Although the discussion of good moral advice in the previous section is consistent with the claim that the success conditions of responsible moral deliberation can be exhaustively explained in terms of what the persons deliberating ought ‘objectively’ (as opposed to ‘subjectively’) to think or do, it does not entail that claim. For discussion, see e.g. Frank Jackson, ‘Decision-theoretic Consequentialism and the Nearest and Dearest Objection’, *Ethics* 101 (1991), pp. 461-482.

¹⁵ For two recent works in which some version of this strategy plays a prominent role, see e.g. Matthew Kramer, *Moral Realism as a Moral Doctrine* (Oxford: Blackwell, 2009) and Ronald Dworkin, *Justice for Hedgehogs* (Cambridge: Harvard University Press, 2011),

that moral thought is universally embroiled in falsehood is a sense in which that claim is of lesser philosophical significance than is sometimes assumed. Yet the discussion in the previous section is consistent with the existence of deep and interesting questions regarding the success conditions of moral claims, such as whether it makes sense to think of ourselves as completely deluded about how we should live, whether moral insight could be completely beyond us, or whether all genuine moral authority obtains relative to some contingent desire, interest or social framework. Nothing said in the previous section should therefore be taken to imply that there is no such thing as a coherent activity of describing the success conditions of moral claims, according to which some moral claims fail this test. The conclusion to draw from the discussion in the previous section is a different one, namely that the only plausible way to carry out such an activity while truly capturing the point of making moral claims is one that takes the activity of moral thinking seriously on its own terms, and therefore one that respects the fact that it is possible for someone to respond more or less successfully to a difficult moral predicament, for example by giving or taking genuinely good moral advice.¹⁶

A different way of describing the position articulated in the previous section is as follows. Suppose our best account of the semantics and ontology of moral thought implies that all moral claims are ‘false’; that the entities they quantify over do not ‘exist’; or that morality is a ‘fiction’. If so, these claims are not best read as implying that moral thought is thereby universally mistaken or defective, or that it should figure on the ‘debit’ column of a ‘cosmological balance sheet’. Instead, these claims should be read as theoretical classification devices (the terminology of which is to some extent arbitrary) that assign at least some moral claims to one rather than another line on the ‘credit’ column of said balance sheet, given a certain convention for classifying different kinds of ‘credit’. In other words, the semantic and ontological projects in question would have to separate the question of whether moral claims are ‘false’ from the question whether there is anything amiss with them.¹⁷ The resulting picture is one on which there are at least two kinds of ‘universal moral error theory’: one which implies that there is something seriously amiss with all moral claims; the other of which does not. Whatever the respective merits of these theories as intellectual constructions, the former kind of theory is as implausible as the latter is toothless.

5. Six objections

There are several obvious objections to the diagnosis given of universal error theories in the previous sections of this paper. Here I shall mention six.

First, it might be objected that everything I have said about the potential for improving our moral beliefs and getting things right in the case of giving and taking good moral advice may well be true and good, but it is only true and good insofar as there are epistemic truths about how we should respond to our moral predicament. It does not follow that there are any moral truths to this effect.

¹⁶ It might be thought that the case of religious belief constitutes a ‘hard case’ analogy for the position adopted with respect to moral thought in this paper. I agree that it does, but will not pursue the consequences of this concession further here.

¹⁷ Arguably, some views of moral thought known under the heading of ‘fictionalism’ are consistent with this interpretation of the success conditions of moral claims. See e.g. M. Kalderon, *Moral Fictionalism* (Oxford: Oxford University Press, 2005).

In response, it might be tempting to invoke ‘companions in guilt’; either by arguing that epistemic truths are no less philosophically mysterious than moral truths or by arguing that, at least in the case of ‘thick’ evaluative terms, epistemic and moral truths are inextricably entwined.¹⁸ Although I have some sympathy for at least the second of these responses, there is no need to invoke either of these strategies here. Even if moral truths are philosophically more problematic than epistemic truths, and even if moral and epistemic truths can be logically prized apart, the question of how to respond to an epistemic predicament involving a moral issue is, at least on some occasions, itself a moral question. The question of how to act in a case of moral uncertainty and epistemic asymmetry is, in part, a moral question about how to act in this kind of situation. In a similar way, the question of how to respond to empirical evidence about the causes of our moral beliefs is itself in part a moral question about how to take responsibility for those beliefs.¹⁹ To think that the truths we need to make sense of good moral advice are purely epistemic, and therefore non-moral, is both descriptively and morally implausible.

Second, it might be objected that everything I have said about the potential for improving our moral beliefs and getting things right in the case of giving and taking good moral advice may all be true and good, but it is only true and good insofar as there are prudential truths, derivable from facts about what is in our self-interest, about how we should respond to our moral predicament in cases of moral uncertainty and epistemic asymmetry. It does not follow that there are any moral truths to this effect.

Once more, in response to this objection there is no need to appeal either to the claim that prudential truths, or facts about self-interest, are as philosophically problematic as moral truths, nor to the claim that at least some prudential truths are inextricably entwined with moral truths.²⁰ Instead, we should deny that all cases of giving or taking good moral advice are better understood and evaluated by interpreting them in purely prudential terms. Not only is this an implausible (and potentially disrespectful) interpretation of what it means for people to give and take good moral advice in at least some actual circumstances. It is also very questionable advice to give to someone experiencing a case of moral uncertainty and epistemic asymmetry regarding how they ought to respond to their predicament, where that is in part a moral question. True, in some cases it can be helpful to illuminate a moral decision problem by modeling it purely in prudential terms. Yet if the practical consequences of doing so have anything to do with it (which they will

¹⁸ This strategy plays a central role in Terence Cuneo, *The Normative Web* (Oxford: Oxford University Press, 2007), and in Hilary Putnam, *Reason Truth and History* (Cambridge: Cambridge University Press, 1981). For wider discussion of the strategy, see Hallvard Lillehammer, *Companions in Guilt: arguments for ethical objectivity* (Basingstoke: Palgrave Macmillan, 2007). What if we propose an error theory for both moral and epistemic claims? From the perspective of the argument in this paper, that would only make things worse. (See e.g. Streumer, ‘Can We Believe the Error Theory?’.)

¹⁹ For a parallel argument in the causal case, see Hallvard Lillehammer, ‘The Epistemology of Ethical Intuitions’, *Philosophy* 86 (2011), pp. 175-200.

²⁰ For a discussion relevant to the latter issue, see e.g. Joseph Raz, *Engaging Reason: On the Theory of Value and Action* (Oxford: Oxford University Press, 1999), pp. 173ff.

on any plausible account of moral deliberation), this will not be true in every case.²¹

Third, it might be objected that the appeal to good moral advice as a problem case for a universal moral error theory is defective because, as described by me, those cases are not strictly speaking ‘moral’. True, people experiencing moral uncertainty have to decide what to do in the context of a particular set of contextual and normative parameters relative to which they can decide more or less responsibly. It does not follow that our grasp of the success conditions that apply to the moral beliefs they form in response to these contextual and normative parameters will survive once these contextual and parameters are lifted. It is one thing to ask what someone ought to do given a set of contingent contextual and normative assumptions. It is quite another thing to ask what someone ought to do independently of any contingent contextual or normative assumptions. It is only the second kind of question that is ‘distinctively’, or ‘substantively’, moral.

In response, assume for the sake of argument that there actually is such a thing as a ‘purely’ moral question that remains once all contingent contextual parameters are lifted.²² If so, the basic conflict between universal moral error theory and the discussion in the previous sections of this paper may boil down to a terminological dispute, but only at the detriment of universal moral error theory. The interest of a metaethical theory depends partly on the extent to which it gives an account of ‘moral’ claims as actually made. If ‘moral’ claims as actually made are claims the correct interpretation of which is that they concern what people ought to do independently of any contingent contextual or normative parameters, then it is an empirically tractable question when, if ever, we either do, or should, make ‘moral’ claims in the required sense. My discussion of the case of moral uncertainty and epistemic asymmetry in the previous section was meant to illustrate that at least some claims actually made in the face of what subjects would themselves describe as ‘moral’ predicaments need not be interpreted as ‘purely’ moral in this sense. On the contrary, at least some claims made in the face of such predicaments quite obviously presuppose a range of contextual and normative parameters, including existing normative commitments that are not all (or at least not all at the same time) in question when making the relevant moral claim. On the required interpretation of ‘moral’, therefore, a universal moral error theory fails to capture at least some recognizably moral claims; in which case it fails to speak to the full range of moral claims; in which case its universality is purchased at the cost of making its truth irrelevant to certain aspects of moral thought.²³

Fourth, it might be objected that everything I have said about the potential for improving our moral beliefs and getting things right in the case of giving and taking good moral

²¹ Here I disagree with the argument in Richard Garner, ‘Abolishing Morality’, in R. Joyce and S. Kirchin (eds.), *A World Without Values*, pp. 217-33.

²² It is not clear that a defender of universal moral error theory is entitled to that assumption. On one (unfriendly) interpretation, it is precisely the non-existence of any coherent question of this kind that (mistakenly) leads the universal error theorist to attribute to moral thought a constitutive commitment to false, incorrect, or otherwise erroneous claims. For a related point, see e.g. Steven Finlay, ‘The Error in the Error Theory’, *Australasian Journal of Philosophy* 86 (2008), 347-69.

²³ It is arguably his acceptance of this fact that allowed Mackie to think that the transition he made half way through his *Ethics* was a theoretically coherent one.

advice may well be true and good even in a moral sense, but it is only true and good in a moral sense ‘within the fiction’ – the fiction of morality.

In response, we should concede that there is more than one coherent way of sorting bits of language and bits of the world such that only some declarative sentences get classified, relative to that way of sorting things, as the genuinely ‘factual’, or ‘truth-stating’ ones.²⁴ What does not follow is that declarative sentences excluded from this class must thereby be described as either ‘false’, ‘erroneous’, or otherwise mistaken. As beings who have the ability to think about the world by means of a plurality of descriptions, we should be able to distinguish the idea that a claim is ‘false’, or ‘incorrect’ on the one hand, from the idea that this claim is not ‘fundamental’ or indispensable for all explanatory (or justificatory) purposes on the other. True, some ‘non-fundamental’ claims about our place in the universe are obviously incorrect. Thus, some people are worse at scrabble than they think they are. Yet not all ‘non-fundamental’ claims are incorrect. Thus, some people are very bad at scrabble. To say that someone is bad at scrabble may be a truth that is in some sense irreducibly ‘perspectival’, in that it obtains only relative to the constraints and parameters of a social practice that involves people adopting the rules of a certain kind of board game.²⁵ It does not follow that being bad at scrabble is something you never ‘really’ are, or that you can only be so ‘within the fiction of scrabble’. In order to know that you are bad at scrabble it is not necessary to be deluded, to speak falsely, or pretend to do anything. Badness at scrabble is literally achievable (and has actually been achieved) by a significant number of real individuals.

Fifth, it might be objected that what I have said about the potential for improving our moral beliefs and getting things right when giving and taking moral advice may well be true and good even in a moral sense, but it is mistaken to think that there is any descriptively adequate way to describe someone’s moral predicament that makes use of normative, evaluative, or moral terms at all (except, that is, as ‘non-factively’ embedded in the content of propositional attitudes).

In response, it might be tempting to get bogged down in the ongoing debate about the ontological credentials of moral explanations.²⁶ Instead, I suggest we make use of the materials employed in the response to the first two objections. The question is whether all cases where good moral advice is offered and taken are better understood by interpreting them in ‘purely descriptive’ terms. This is partly a question of how it is best to think about such scenarios, and the way people do, and ought to, respond to them. In order to answer this question, it is not necessary to deny that there are purely descriptive aspects of every conceivable moral or epistemic predicament (whether ‘physical’, ‘biological’,

²⁴ Perhaps it could be helpful to say this about declarative sentences the correct use of which exhibits a causal relationship to what they are about, or that play some privileged role in the explanation of a given set of data, such as experimental observations under controlled conditions, or the experience of unmediated sense data. For similar observations, see e.g. Hilary Putnam, *Ethics without Ontology* (Cambridge: Harvard University Press, 2004) and Huw Price, *Naturalism without Mirrors* (Oxford: Oxford University Press, 2011).

²⁵ For a systematic exploration of the idea that some truths are irreducibly ‘perspectival’, see e.g. Williams, *Ethics*, and A. W. Moore, *Points of View* (Oxford: Oxford University Press, 1997).

²⁶ See e.g. Nicholas Sturgeon ‘Moral Explanations Defended’, in J. Dreier (ed.), *Contemporary Debates in Moral Theory* (Oxford: Blackwell, 2006), pp. 241-62.

‘economic’, ‘material’, ‘psychological’, or ‘decision theoretic’) that ought to play a significant part in our interpretation of those predicaments.²⁷ The question is why, in compiling the list of legitimate descriptions, we should exclude the moral descriptions in terms of which those people themselves actually resolve to make sense of their predicament as morally engaged subjects. While there is a vast number of possible interpretive or explanatory projects relative to which it makes perfect sense to bracket all morally engaged descriptions of a moral predicament, it does not follow that no interpretive project that employs at least some morally engaged descriptions is able to make genuine sense of that predicament in the way that matters. Of course, if every morally engaged description of a given situation were inconsistent with some true non-moral description of the case, then we would have to admit that there is something amiss with the former on pains of inconsistency. Yet the mere fact that the situation in question has a set of true non-moral descriptions is clearly insufficient to establish that claim.

Sixth, it might be objected that what I have said about the potential for improving our moral beliefs by giving and taking good moral advice is all true and good, but it actually presupposes exactly the kind of semantic and ontological theory (such as a theory of the supervenience relation between moral and non-moral truths, facts and properties) that I have been failing to engage with all throughout this paper. It follows that my strategy of attempting to sideline the projects of semantic and ontological analysis for the purposes of interpreting cases of moral uncertainty and epistemic asymmetry is inconsistent at the core.

In response, I agree with the basic thought behind this objection, but deny that it presents a problem for the discussion in the previous sections of this paper. We should accept the claim that there is a genuine question about what earns a given domain of claims the entitlement to be called ‘true’, or in the business of picking out genuine ‘facts’ or ‘properties’, and so on. We should also accept that insofar as we are committed both to moral truths and to a range of descriptive truths on which they somehow ‘depend’, there is a genuine question of how to accurately characterize the relevant dependency relations. What we should not accept is that in earning this entitlement, or in giving the relevant characterizations, our theory of moral thought should be thought of as standing in a relationship of asymmetric dependence on a theory of ‘truth’, ‘property’, ‘fact’ or ‘supervenience’, such that we could end up having to say that genuinely insightful moral claims are actually mistaken or defective, or that hard earned insights into social reality are merely illusory or fictional (in a pejorative sense). Quite how this feat is to be achieved is an interesting challenge, and far from trivial. But it is a challenge for philosophy, not for the credentials of moral thought as such.

6. A distinction without a difference?

Throughout this paper I have assumed that one plausible criterion of a successful metaethical theory is that it make interpretive sense of moral thought as embedded in a social world in which morally serious people describe and communicate their predicament in moral terms, and sometimes respond to that predicament in morally insightful ways. I have argued that no universal error theory is likely to satisfy this

²⁷ For doubts about even this claim, see e.g. Putnam, *Reason, Truth and History*.

criterion because, although it is obviously consistent to say (in one voice) that there are more or less insightful ways for people to think about their moral predicament, and then to say (in another voice) that this is only the case within the 'fiction' of morality, there is no way for a universal error theory to give plausible content to the second claim without deflating it to the point where it becomes an issue of little more than terminological interest. What universal error theories offer is therefore something that comes dangerously close to a 'distinction without a difference': perhaps a plausible way to distinguish moral thought with respect to other forms of thought from which it is in some (potentially interesting) ways different, but not a plausible way to make sense of the idea that all moral claims are mistaken, defective, or somehow based on an illusion. Once we accept the idea that there are people who are right to take their moral predicament seriously on its own terms (such as the case where someone is fortunate enough to receive some good moral advice), there is no credible interpretation of what they are doing according to which every time they sincerely utter a declarative moral sentence, they are thereby embroiled in some kind of philosophically diagnosable error or mistake.

*Churchill College,
Cambridge CB3 0DS
hallvardlillehammer70@gmail.com*