Ethics vs. Metaphysics

Michele Odisseas Impagnatiello

(Forthcoming in *The Journal of Philosophy*)

Abstract

Sometimes, a metaphysical theory has revisionary ethical consequences: for example, some have thought that modal realism entails that there are no moral obligations. In these cases, one may be tempted to reject the metaphysical theory on the grounds that it conflicts with commonsensical ethics. This is an ethics-to-metaphysics inference. My claim is that this inference is in general irrational, and that the fact that a metaphysical theory has highly revisionary ethical consequences is no reason at all to reject the theory. I argue for this claim on the basis of general epistemic principles about the transmission of justification, and what makes for a good argument. Furthermore, I argue that my account can explain why a certain narrow class of ethics-to-metaphysics inferences are rational.

1 Introduction: Ethics meets Metaphysics

Metaphysics is about the way things are: what there is, what it's like, the nature of things, et cetera. Ethics is normative: it's about what we should and shouldn't do, what's good and bad, et cetera. There are connections between the two: what we should do partly depends on how things are, after all. I shouldn't punch Jimmy because that would cause him pain. But some of these connections appear to be far reaching, and some metaphysical theories appear to have highly revisionary ethical consequences. For example, some have thought that if modal realism is true, then we don't have any moral obligations;[1] others have claimed that for there to be moral truths, God must exist.[2]

In these cases, one might be tempted to draw an inference from ethical premises to metaphysical conclusions. Surely, we have some moral obligations. If

[1] Adams (1974); Heller (2003).
[2] See Baggett & Walls (2019) for a history of this kind of argument for God's existence.

modal realism is incompatible with that, then so much the worse for modal realism. In general, such an argument has the following form:

P1 [ethical claim];
P2 if [metaphysical theory], then not-[ethical claim]; therefore,
C not-[metaphysical theory].

Is such an argument any good? If you think so, then you're a *Moralist*. You think that ethics can dictate metaphysics, or at least it can give us some reason for or against a metaphysical theory. If you think such an inference is not good, then you're an *Anti-Moralist*. You think that we cannot use ethical premises in arguing for or against a metaphysical conclusion.

Which view is right? You may think there is no general answer, and that it really depends on the subtleties of each particular case. Or you may think that the answer depends on one's prior methodological standpoint: some people might give more weight to ethics, and be happy to let ethics dictate metaphysics, while others might prefer to leave metaphysics alone.

In this paper, I'll argue that the situation is less flexible than this, and that there is a principled answer as to whether Moralism or Anti-Moralism is true. The answer depends on the structure of the ethics to metaphysics inference in question. For the main class of such cases, I'll argue that Anti-Moralism is surprisingly right. In some limited cases, Moralism may be right.

In section 2 I put forward the main kind of ethics to metaphysics inference. In sections 3, 4, and 5 I argue that, surprisingly, Anti-Moralism is the right account of those inferences. In sections 6 and 7 I consider two different kinds of ethics to metaphysics inferences, where Moralism might be the right response.

2 Commonsense Ethics against Revisionary Metaphysics

Metaphysicians say wild things at times. Some metaphysicians believe that there are no people or tables or chairs, or that there is an object composed of a person plus the Eiffel tower, or that all possible worlds are equally real, or that nothing is conscious, or that everything is conscious, or that near me there's a myriad of people-like objects, slightly different from me. Normally, that doesn't bother us very much.

However, in some cases, these wild metaphysical theories have consequences for things we care about.

For example, consider modal realism. That's the view, famously defended by Lewis (1986), according to which all possible worlds exist and are just as real and concrete as our world. So, just as you and I exist, there are flying pigs and purple cows; they're just spatiotemporally disconnected from us. Never mind whether this view is true. Suppose it is true, as Lewis thought. What of it?

Adams (1974) pointed out that modal realism appears to have consequences for things we care about. For example, we typically only care about actual sufferings, and not about merely possible sufferings. Joe is happy; the fact that Joe *could* be in great pain doesn't worry us, given that he's happy. And the reason we don't care about possible sufferings is that they're not *real*. But if modal realism is true, then there is a concrete possible world, just as real as this one, in which a counterpart of Joe *is* in great pain. This suffering Joe is not an element of our spacetime, but intuitively, that isn't morally relevant. It's *as if* Joe is happy in this room, but in the next room over a clone of Joe is in pain. Since we care about people suffering, modal realism does matter for things we care about.

This clearly generalizes. For all the ways things could be that we care about, there is a way like that, both in the good and the bad. There are communities of perfectly happy people, and communities of people tortured forever. All of it is real; it's just the "next room over". This already is a bit unsettling; a sense of moral vertigo enters into me when I consider that.

But the problem trickles down to more practical matters – to actual decisions, as Heller (2003) argued. To take the standard ethics cliché, suppose I'm strolling by a lake, and I see a drowning child crying for help. Now, we might disagree about morality – consequentialism, deontology and all that – but I clearly have to jump in and save the child. It might ruin my clothes, but I have to do it. For simplicity, we can assume that there are only two possibilities: I jump in and the child lives, or I do nothing and the child dies. Normally, I would try to actualize the first one, because it's the better one. But if modal realism is true, both these possibilities are real. In one world, I jump in and save the child; in another, I don't and the child dies. No matter what, one child dies and another one doesn't. What's the point of saving this child? I'm not making reality a better place. I'm just choosing whether this spatiotemporal mass is the one with the child who drowns or the child who lives. In

a way, it's as if there are two towns I could live in. In one of them, a child lives; in the other, a child drowns. Do I really have an obligation to live in the one where the child lives? It doesn't seem so.

If this is right, then modal realism has the consequence that, contrary to what we thought, I actually don't have to save the drowning child. And so, one could give the following argument:

P1 Modal realism is true;
P2 If modal realism is true, I don't have to save the drowning child; therefore,
C I don't have to save the drowning child.

The conclusion is very surprising, for it seems pretty obvious that we have to save the drowning child. Not only surprising, but unsettling and disturbing. It seems like metaphysics can teach us a whole lot about our ethical obligations, and radically alter our view of what we should do: after all, the argument clearly generalizes. Not only do I not have to save the drowning child, but I don't have to help anyone in need, and it's not wrong to murder or torture, et cetera.

Of course, this argument isn't airtight. The reasoning sketched above isn't flawless, and had a consequentialist flavor. Maybe it really matters whether *we* cause good outcomes, even if the overall consequences are the same; maybe people who are spatiotemporally related to us matter more than people who aren't. In other words, maybe P2 isn't true. Maybe. Still, the reasoning is plausible, and has at least some credibility. For the argument's sake, let's treat P2 as true.

This argument here is an argument from a metaphysical premise to a surprising ethical conclusion. But rather than surprising or unpleasant, one might think the conclusion is just false. After all, someone's modus ponens is someone else's modus tollens. If you don't want to accept the conclusion, all you need to do is "flip" the argument, as follows:

P1 I have to save the drowning child;
P2 If modal realism is true, I don't have to save the drowning child; therefore,
C Modal realism is false.

Now we have an argument against modal realism, with ethical premises. This is an example of inference from ethics to metaphysics, as exposed in the previous section.

Let's call this kind of argument the "Moralist Argument". There are other cases like this one, but for simplicity I will just focus on this example.[3]

Is it a good argument? Let me first be clear on what this question is asking. Sometimes, in philosophy a good argument is understood as one that is logically valid and sound. The moralist argument above is clearly valid, and if modal realism is false it's sound. But there is another sense of good argument. That's the sense in which an argument is *successful*, in that it can rationally convince someone of its conclusion. Another way of saying it is that a good argument is one that gives us reason to believe the conclusion on the basis of the premises. An argument being valid and sound isn't enough for it to be good in this sense. For example, the argument "*p*; therefore *p*" is valid and sound if *p* is true, but doesn't give us reason to believe *p*. As we'll see, arguments aren't bad or good simpliciter; they're only good or bad relative to people in certain epistemic positions. In the rest of the paper, by "good argument" I mean one that is good in this sense. This is the sense that is relevant, because we're wondering whether we gain reason to reject a metaphysical theory on the basis of ethical premises.

Moralists think that this argument is a good argument. It gives us at least *some* reason to disbelieve modal realism. Some Moralists might be punchier, and would say that this argument basically defeats modal realism in and of itself. Anti-Moralists, in contrast, think that this argument gives us *no reason at all* against modal realism. According to them, in whatever high or low regard we held modal realism, we shouldn't change our mind at all given the above argument.

Which is the right view? At first glance, it's hard to resist the Moralist temptation. The argument above is valid after all, and presumably sound as well. Surely, the fact that modal realism has this strange consequence should tell against modal realism at least to some degree! As long as it's *a* reason, no matter how weak, against modal realism, then Moralism is vindicated. And surely, we do gain some reason. Isn't it obvious that we should save the child? An uncompromising Moorean will say that it's so obvious, that modal realism just goes out of the window. But even without going that far, it's plausible that in our theorizing we should balance

[3] Another widely discussed case is the *personite problem*. Johnston (2016) has argued that if four-dimensionalism is true, then it's not morally permissible to, say, study Hungarian. The Moralist argument is then to infer that four-dimensionalism is false, from the premise that it's permissible to study Hungarian. I won't discuss the personite problem in this paper, but the conclusions I reach apply to this Moralist argument as well.

all the relevant considerations. It seems that we should save the child, and this is going to exert some pull away from modal realism. Lewis himself though that if modal realism has strange ethical consequences, then this is some reason to disbelieve it (his own solution was to deny that modal realism has such consequences: he denied P2 in the Moralist argument).

I will argue, however, that we have to resist the Moralist temptation. Anti-Moralism is right: the argument above gives us no reason at all against modal realism. This might seem surprising. As a warm-up, let me put forward a similar inference, from ethical premises to a different descriptive conclusion, in which Anti-Moralism is clearly the right answer.

Take the drowning child again. We have to save the child; that's true if anything is. But even putting modal realism aside, there's a variety of non-metaphysical hypotheses about the way the world is that, if true, would contradict that claim. Suppose that at the other end of the lake a hundred children were drowning (and could be saved by the pull of a lever, say). Then we don't have to save this child; rather, we have to run to the other end of the lake and save the one hundred children before it's too late. Or suppose that this isn't actually a child, but a childlike doll. Then we wouldn't have to save it; dolls don't matter that much.

But clearly, we have to save the child! So, here's the Lake argument:

P1 I have to save the child;
P2 If there are one hundred drowning children over there, I don't have to save the child; therefore,
C There aren't one hundred children drowning at the other end of the lake.

I trust that this argument isn't good. It's valid, and it may be sound. But something seems off. We can't really, on ethical grounds, conclude that the lake doesn't contain drowning children over there. The hypothesis that one hundred children are drowning over there has a strange ethical consequence (that I don't have to save this child); but the fact that it has a strange ethical consequence is *no reason at all* against the hypothesis. Anti-Moralism is clearly true of this and of similar arguments constructed using other empirical hypotheses.[4] For these reasons, let's call these the "easy cases" of Anti-Moralism.

[4] For another example, consider the Doll argument:
P1 It's wrong to sit down and do nothing;

Of course, these arguments are different from the modal realism one. The conclusion here isn't a metaphysical claim, but a perfectly ordinary, empirical claim about the distribution of drowning children. This might make a big difference. As it happens, it doesn't. Indeed, we certainly need an explanation of why Anti-Moralism is intuitively right in the Lake case. I'll provide such an explanation. Once the explanation is given, we will see how it also applies to the original argument targeting modal realism. The Lake argument is a bad argument in that it doesn't give us reasons to believe the conclusion. I'll argue that this is because those two arguments *fail to transmit justification* from the premises to the conclusion. To see why, we need to take a detour into epistemology.

3 Transmission Failures and Inferential Justification

Let's leave metaphysics and ethics behind, just for a section. And let's look at another argument, which, I hope, will prove instructive with respect to the Moralist arguments we've been considering.

I live in a pretty modest lodging; I'd prefer to live in a giant mansion. Will I ever? Maybe some day. But in the short term, I'm pretty sure I won't. In other words, I'm pretty sure that next month I won't live in a giant mansion. That's something I justifiably believe (arguably, it's something I know). Now suppose that we're talking about my finances. And somehow, the question comes up: will I become rich in the next month? I say no. You ask me for a reason. I oblige, and give this argument:

P1 I won't live in a mansion anytime soon;
P2 If I become rich next month, I will live in a mansion soon enough; therefore,
C I won't become rich next month.

Let's call this the Rich argument.[5] There is a sense in which the argument is good. P1 is very reasonable. P2 is reasonable as well: I'd love to live in a giant mansion, so if I do become rich I'll certainly buy one. And C follows from P1 and P2, from a

---

P2 If this is a childlike doll, it's not wrong to sit down and do nothing; therefore,
C This is not a childlike doll.
Which intuitively is just as bad as the Lake argument.

[5] Arguments similar to this one were examined by Vogel (1990) as possible counterexamples to the principle of knowledge closure.

simple modus ponens. So we have a valid argument (most likely a sound one), with premises which are justified and even known!

Still, it's pretty clear that something is wrong with the Rich argument. Suppose someone raised the possibility that I *will* become rich, if, for example, some random stranger decides to buy my hat for a few million dollars. I can't rule that possibility out, or even reduce my confidence in it by any amount, using this argument. I can't really give, as a reason to believe that I won't become rich, that I won't live in a giant mansion.

Why is that? Well, one of premises of the argument is that I won't live in a giant mansion. That's a justified belief, to be sure. But what justifies it? What justifies it is, I think, something like the following two beliefs: first, that I won't become rich anytime soon. And second, that if I won't become rich anytime soon, I won't have a mansion anytime soon. Both these beliefs are justified: why would I become rich soon? I have a low paying job, and no great prospect in sight. And I know that the most likely way to get a giant mansion is by being rich; it's not like many non-rich people somehow have giant mansions.

But then it's pretty easy to identify what's wrong with the Rich argument. What's wrong is that I believe I won't live in a mansion *because* I believe I won't become rich soon. To be more precise, I'm justified in thinking that I won't live in a mansion on the basis of my antecedent justification in thinking that I won't get rich. I can't then justify the latter on the grounds of the former. To do so would be epistemically circular (at least, in ordinary circumstances. We can imagine exotic cases in which this is not the case, involving crystal balls or time travelers. I will address this later).

Here, I'm appealing to the notion of *epistemic priority* or *epistemic dependence*. While no account of epistemic priority is widely agreed upon, the notion is an intuitive one. Some of our justification is dependent on other justification. In the case at hand, my justification for the proposition that I won't live in a mansion is *dependent*, or *inferential*, on my justification for the proposition that I won't be rich. The notion of justification at play here is *propositional justification*. This is the kind of justification that a proposition has for a subject when the subject is justified in believing it (if, for example, the subject has sufficient evidence to believe it), whether or not they believe the proposition. This is in contrast with *doxastic justification*, which is the justification that the subject's belief has (if, for example,

it's formed on the basis of the right evidence).[6] Talk of "inferential justification" may suggest a doxastic reading, where a belief is formed inferentially. Here, and in what follows, I use "inferential" to apply to propositional justification, when the justification rests on other justification.

We can see the problem if we ask for the larger argument that makes clear the source of out justification for the first premise:

P0 I won't become rich next week;
P0.5 If I don't become rich next week, then I won't live in a mansion; therefore,
P1 I won't live in a mansion;
P2 If I become rich next week, then I will live in a mansion; therefore,
C I won't become rich next week.

Where we can directly see that the conclusion of the argument is one of the premises.

This is an example of a *transmission failure*. The justification that I have for P1 and P2 doesn't transmit to C, even if P1 and P2 logically entail C. Even if the Rich Argument is valid and sound, it fails to transmit justification from the premises to the conclusion. The issue of transmission failure has received a lot of attention in epistemology.[7] Sometimes, it's controversial whether an argument fails to transmit justification: for example, it's an open question whether Moore's proof of an external world transmits justification (see e.g. Pryor (2004) and White (2006)). It should be uncontroversial, however, that there is transmission failure in the Rich argument.

It might seem strange that the Rich argument doesn't transmit justification. After all, am I not justified in thinking that I won't become rich? And isn't that the conclusion of the argument? Yes. I *am* justified in thinking that I won't live in a mansion, and I *am* justified in thinking that I won't be rich. But I am *not* justified in thinking that I won't be rich based on the fact that I won't live in a mansion. It's the other way around: I'm justified in thinking that I won't live in a mansion in virtue of my prior justification that I won't get rich. I was *already* justified in thinking that I won't get rich.

[6] On propositional vs doxastic justification, see Silva & Oliveira (forthcoming).
[7] See, for instance, Wright (1985) and Davies (1998). For overviews on transmission and transmission failure, see Tucker (2010) and Moretti & Piazza (2023).

Crucially, the Rich argument is bad because I don't have any *independent* justification for P1. If I did, then the argument would be fine. For example, suppose I have a crystal ball that allows me to directly see the future. I look into it, and I see myself living in a non-mansion. Now intuitively this is a good reasoning: I won't live in a mansion, therefore I won't get rich. Now I *do* get reason to believe I won't get rich. Now the Rich Argument works. And it does because the crystal ball gives us some independent justification for P1. This by itself shows that whether the Rich Argument is good or bad doesn't just depend on its formal features, i.e. on its structure, nor on what propositions figure in the argument. Structurally, it's a modus tollens – it doesn't get any better than that. But the very same modus tollens is bad in the original scenario, and great if we have a crystal ball.

In the Rich argument, once P0 and P0.5 are added, the epistemic circularity becomes transparent: I believe P1 on the explicit basis of C. But an argument can be similarly problematic without this explicit circularity. Consider the Lottery Argument:[8]

P1 I won't be rich next week;
P2 If I win the lottery tomorrow, I will be rich next week; therefore,
C I won't win the lottery tomorrow.

I trust that this argument feels intuitively just as bad as the Rich argument. However, the same diagnosis can't exactly apply, since C isn't what directly justifies P1. If asked for justification that I won't be rich next week, I might not necessarily say that it's because I won't win the lottery. Luckily, the diagnosis can be extended to account for the badness of this argument.

What is it that justifies my belief that I won't be rich? If someone asked me why I thought that, I would probably say that it's because I have a low paying job. In other words, I cite a cause of my future economic situation, and I am justified in believing that that cause holds. But there are many conditions that are required for me not being rich in the future, in addition to me having a low paying job. One of the conditions is that I won't unexpectedly get a large inheritance from a distant uncle. Yet another is that a big pile of money won't magically materialize into my room. And another one is that I won't win the lottery.

[8] Hawthorne (2004) consider such an argument, in the context of knowledge closure.

In order to be justified in thinking that I won't be rich, I must be justified in all of these conditions. If I wasn't justified in thinking that money won't magically appear in my room, then I wouldn't be justified in thinking that I won't be rich. Of course, such faraway conditions, such as money not appearing out of nowhere, or a distant relative giving me a large inheritance, needn't be explicitly present as such in my thinking (maybe they're all filed under a big "no strange thing will happen" proposition. And I am justified in that catchall proposition, for strange things are unlikely). This isn't a problem, since propositional justification is what is at stake. In order to be justified in thinking that I won't be rich, I must have justification, implicitly or explicitly, for believing in all the conditions.

This crucially depends on the fact that our knowledge of the future is fully inferential on our knowledge of the present (and of the past) and of how it will bring about the future. The only way we could be justified in what could happen tomorrow is by being justified in what happens now (or in the past), and in how it brings about what happens tomorrow. The situation can be represented in figure 1.

Figure 1

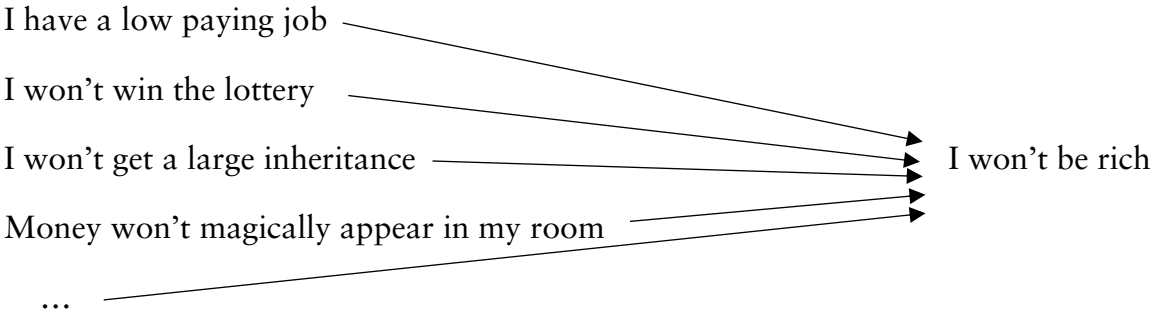

Figure 1. The arrows connect the causally relevant factors to their effect

In figure 1 we can see (some of) the causally relevant factors of my not being rich. And, when it comes to my not being rich in the future, I am only justified in believing it by means of having justification in believing its causes.

This, then, is what's wrong with the Lottery argument. We are justified in believing "I won't be rich" wholly inferentially, by being justified in believing propositions about the causes of the fact that I won't be rich (explicitly or implicitly); we cannot, then, infer back from it to any one of its causes. To do so would be implicitly circular. In general, we can give the following criterion of transmission failure:

> **Transmission Failure.** If our justification for believing a proposition P is fully inferential on its conditions C1, …, Cn, then we cannot infer any of its conditions from P. An argument to that effect would fail to transmit justification.

A crucial requirement of this account of the failure of the Rich and Lottery argument is that the justification for P is fully inferential on P's conditions.[9] There are two ways in which this criterion could not be met, and in those cases the argument would not fail to transmit justification.

The first way is that we could have *direct* justification in P. In the Lottery case, it's hard to imagine having direct justification concerning facts about the future. But let's imagine that time has passed, so that the question is whether I have won the lottery last month. And suppose that I have amnesia and suddenly lost all my memories of the past. I can still be directly justified in believing that I'm not rich, by looking at my bank account and my living situation. Then, it would be rational to infer that I haven't won the lottery in the past (and that money didn't magically materialize, etc.).

Secondly, we could have inferential justification that doesn't wholly come from P's conditions. Suppose a time traveler from the future pops in and tells me that, indeed, I won't be rich. Then it seems that now I can infer that I won't win the lottery next week (and the other conditions). This is because the time traveler's testimony provides me with independent justification that I won't be rich, one that is not inferential on its causes.

As a matter of fact, our knowledge of the future is always mediated by its causes, and that's why arguments like the Lottery naturally strike us as bad. But, as those two variations show, they needn't be: it all depends on how we are justified in the relevant premises. And my account of transmission failure explains why in the time-traveler and amnesia case the argument is good, while in the normal case, it's bad.

---

[9] What's a "condition"? In the cases seen so far, conditions are causes or causally relevant factors. But transmission failures occur more generally, when conditions are the reasons or grounds that make a fact hold. For example, if my justification for believing a conjunction "A and B and … and Z" is fully inferential on the individual conjuncts, an inference from the conjunction to an individual conjunct would intuitively fail to transmit justification, and the criterion I gave accounts for this.

There are a lot of interesting issues that arise here. A full theory of the transmission of justification is certainly needed to understand all that's going on. But for present purposes, a fragment of that theory will be enough. We have enough to understand why the Rich Argument and the Lottery Argument are bad arguments. I will argue in the next section that we can apply the same analysis to the various Moralist arguments I've put forward.

4 Transmission Failures and Anti-Moralism

Let's begin with one of the "easy cases" in which anti-Moralism is clearly the right stance. Here's the Lake argument again:

P1 I have to save the child;
P2 If there are one hundred drowning children over there, I don't have to save the child; therefore,
C There aren't one hundred children drowning at the other end of the lake.

It's pretty obvious that I should save the child. But what justifies me in thinking that?

The conditions for this claim are the descriptive facts of the case, and at least partly, my justification is posterior to the justification that I have for believing the descriptive facts. Some of these descriptive facts readily come to mind. That a child is drowning, for example. I'm justified in thinking that because I see him. And if I didn't think that, I wouldn't think I'd have to save him: if the child was happily swimming, I wouldn't think there's anything I had to do. Furthermore, I think that that I'm able to jump in and save the child. If I didn't think that, I wouldn't think that I have to save him.

How exactly do these two claims provide justification for the claim that I have to save the child? There are two justificatory models of the situation. On the first one, there are overarching moral principles, and these descriptive claims first justify the intermediary proposition that the conditions for the moral principles are met. For example, a consequentialist will say that these propositions justify the intermediary proposition that saving this child has the best consequences, and in turn this proposition justifies the proposition we have to save the child.[10] We may

[10] Different moral theories will plug in different intermediary propositions.

call this the "generalist" model. In contrast, someone might deny that moral principles play a role in justifying the claim that we have to save the child. This person, which we can call the "particularist", will say that these descriptive propositions directly justify the belief that we have to save the child.[11] Either way, the justification that we have for believing these two descriptive propositions is prior to the one we have for believing that we have to save the child.

What about the proposition that appears in the Lake argument, namely that there aren't one hundred children drowning at the other end of the lake? This is another descriptive claim that is relevant to the question of whether we have to save the child. And indeed, it's another condition for it to be true. On the generalist model, this claim is another condition for the intermediary proposition. Using consequentialism as an example, that the lake doesn't have a hundred children drowning over there is a condition for it to be the case that saving this child has the best consequences (if there were those children, then saving those children instead would have the best consequences). On the particularist model, this claim is another direct condition for the claim that we have to save the child (otherwise, it would be a mystery why it's relevant to it). On either model, this claim is epistemically just like the claim that a child is drowning: it's prior to the claim the we have to save the child. We're justified in thinking that we have to save the child on the basis of this claim (among others).

Of course, in this case our justification is implicit. We don't explicitly give, as justification for the claim that we have to save the child, that the lake doesn't contain one hundred children drowning at the other end. But that doesn't matter; we still must be justified in believing it. If we weren't (if, say, we heard some rumors of such hundred drowning children) we wouldn't be justified in thinking that we have to save this child.

We can now see why the Lake argument is a bad argument: like the Rich argument, it fails to transmit justification.[12] We can't infer, from the claim that I have to save the child, that there aren't one hundred children drowning at the other

[11] In the literature, generalism and particularism are theses about the role that moral principles play in moral theory more generally. On that, see Ridge & McKeever (2016). Here, I'm using those labels only to refer to the role that moral principles play in the structure of our justification for believing moral facts. Someone could be a particularist in this sense even if they believe that there are general moral principles, as long as our justification doesn't go through them.

[12] The same diagnosis can be applied to the Doll argument (see footnote 4), since "this is not a childlike doll" is similarly a condition for the claim that I have to save the child.

end of the lake, because the latter is a condition for the former,[13] and our justification for the former is wholly posterior to (or inferential on) the justification we have for its conditions. Once we're satisfied with this explanation, it's an easy step to conclude that the original Moralist Argument about modal realism is bad as well. Here's the argument:

P1 I have to save this child;
P2 If modal realism is true, I don't have to save this child; therefore,
C Modal realism is false.

The falsity of modal realism is just another condition for the claim that I have to save this child. For example: for the consequentialist, that there aren't two equally real possible worlds, one in which I save the child and one in which I don't save him, is a condition for it to be true that saving this child has the best consequences. In general, inasmuch as we take Modal realism to be relevant to the question of whether or not we should save the child, either it's a direct condition of it, or it's relevant because by being relevant to the purely descriptive facts about the lives and death of children, in the same way that the hypothesis about the lake containing one hundred drowning children is.

In the justificatory structure of the situation, the falsity of modal realism is prior to the mixed ethical claim. As before, this justification is not present explicitly in our reasoning, but it doesn't matter. We must have justification for thinking that Modal realism is false, for us to be justified in thinking that we have to save the child. If we weren't (suppose that the arguments for modal realism were convincing), then we wouldn't be justified in thinking that I have to save this child. And we *are* justified in thinking that modal realism is false. It's an extremely revisionary view, after all. This is why the Moralist Argument is bad: we can't infer a condition from something that we're only justified in believing inferentially, from its conditions. The Moralist argument is a case of transmission failure.

This might be surprising. Isn't it obvious that I should save this child? Yes, it is. But it's obvious only because it's antecedently obvious, among other things, that there aren't one hundred children drowning over there (why on earth would that be?), and that this is indeed a child and not a doll (dolls so realistic are very rare), and that modal realism is false (that's a wild metaphysic). The obviousness of this

[13] In this case, the descriptive facts are conditions in virtue of being grounds, rather than causes (see footnote 9).

mixed ethical claim is parasitic on the obviousness of descriptive facts, empirical and metaphysical alike, that are the conditions for this fact to hold.

The diagnosis can be suitably adapted to all Moralist arguments of this form, namely those in which a mixed ethical fact is the premise.[14] In general, such mixed ethical facts are justified inferentially, from the various descriptive facts of the case, including which metaphysical theory is true. For this reason, one can't reject the metaphysical theory on the basis of this ethical fact, and such an argument doesn't transmit justification. When it comes to this kind of argument, then, Anti-Moralism is true in general.

In the next section, I consider an objection to my argument, by way of rejecting the structure of the ethical justification that I've put forward.

5 Ethical Justification and Moral Perception

My argument for Anti-Moralism crucially rests on the claim that our justification for the mixed ethical claim is inferential from the descriptive facts of the case. It's tempting to resist this picture of ethical justification. One could claim that we have *direct* justification in ethical claims. If that were true, then the arguments wouldn't fail to transmit justification. Rather, it would be like the Lottery argument in the case where I directly observe that I don't live in a mansion.

And indeed, that is a tempting thought. It does seem that when we see the drowning child, we directly get justification that we ought to save it. There's no reasoning involved, no sophisticated deduction. It's just direct intuition. We *feel* that we have to save the child. Some would say that we *see* the moral fact that we have to save the child. They say we have moral perception, that gives us direct justification in normative claims, just like ordinary perception gives us justification in descriptive claims.[15]

[14] To illustrate: in the case of personites and four-dimensionalism, our justification to believe that it's permissible to study Hungarian is posterior to our justification to believe that (for example) studying Hungarian benefits me and harms nobody; and in turn our justification to believe this claim is posterior to our justification to believe that four-dimensionalism is false.

[15] For a defense of moral perception, see McGrath (2004). For an overview of the debate, see Werner (2020).

Whether moral perception is inconsistent with my diagnosis depends on what moral perception is taken to be. Moral perception could just be the claim that our belief that we have to save the child isn't formed inferentially from prior descriptive beliefs, but rather is formed directly upon seeing the drowning child. But this isn't inconsistent with my account. What matters is whether our *justification* for the belief is epistemically posterior to our justification in the descriptive claim that a child is drowning, etcetera. How the belief is actually formed is independent of that question.

A second way moral perception may seem relevant is that it also doesn't seem that our *justification* is inferential. In particular, the philosopher sympathetic to moral perception may say that we cannot really articulate the descriptive and normative conditions that justify our belief that we have to save the child. But again, I can concede that. Whether our justification is epistemically posterior doesn't depend on whether it's explicitly present in the head, because what matters is propositional justification and not doxastic justification. This is true as it can be seen from the case of the Rich argument. I cannot really articulate all the epistemically prior propositions that justify my belief that I won't be rich. But it's still true that it's inferential on all claims such as the claim that I won't win the lottery, as well as claims I never explicitly considered such as the claim that money won't magically appear in my room, or that a stranger won't buy my hat for a billion dollars, and so on. Since that move is right in the case of a propositions like the one that I won't be rich, it can also be right in this case.

Finally, moral perception can be taken as the claim that perception directly justifies us in the mixed ethical claim that we have to save the child. Let's call this view "naïve moral perception". This view *is* a challenge to my view. However, I think that this view is implausible. Firstly, the intuition that our justification is not inferential can be put under pressure. Say we're walking by the lake, we see the drowning child, and you say "we have to save him". I ask you why you believe that. At the beginning, you might say "isn't it obvious? Can't you see?". But if I nudge you, you would eventually say something like "well, the child is drowning!". Clearly, we think we should save the child because we think the child is drowning. We surely don't think the child is drowning because we think we should save him: that would be to take things backwards. This is reason to believe that the justification is actually inferential and not direct, even if at first glance it may appear direct.

Secondly, and more importantly, naïve moral perception has weird consequences. If we gain direct justification that we have to save the child, then Moralism is right about the easy cases too: we can gain justification for the claim that there aren't one hundred children drowning at the other end of the lake, on the basis of the direct moral justification that we have to save this child. But this is absurd. The moral sense may be strong, but it cannot tell us what's at the other end of the lake; it's not that strong. That's something we can only find out by going there and checking, or through background information about the distribution of drowning children in lakes. A way to see why this is implausible is to imagine that we're just wondering for the sake of wondering, before anything dramatic happens, whether there are one hundred children drowning over there. Then we see this drowning child, and we see that we ought to save him. It seems very strange that we suddenly gain extra justification that there aren't one hundred children drowning at the other end of the lake.

So naïve moral perception must be rejected. This isn't to say that the whole theory of moral perception has to be rejected. As said earlier, I can concede that we form the moral belief without doing any inference; my claim is about the structure of justification. Furthermore, we can find an alternative sense in which there *is* direct moral perception. Rather than say that moral perception gives us direct justification in the mixed moral claim, we can say that moral perception gives us direct justification in the particular moral conditional: in the claim that *if* things are as they appear (a child is drowning, no one else is, etc.), then I should save the child. After all, I would have the same moral perception if the whole set-up was an illusion, and there was no child and no lake. It would still be true that *if* things are as they appear, I have to save the child. In this sense, moral perception is direct in that we have direct knowledge of the particular moral conditional fact of the case. This is to be contrasted with an intellectualist view on which in each case we get to the mixed moral fact by applying some overarching moral principles (e.g. "maximize total happiness") to the descriptive facts of the case at hand.

All in all, our justification for the first-order ethical claim is inferential, and so the Moralist argument is a case of transmission failure. However, it's interesting to consider whether it could ever be a good argument. After all, the Lottery argument is a bad argument only in the standard situation in which we know the future through the present. If a time-traveler told us we won't live in a mansion, then we can infer we won't live the lottery. What would be the analogous case for the

Moralist argument? It would be a case in which our justification for the first-order ethical claim is either direct, or gained inferentially without passing by the descriptive facts and the relevant moral law. This never happens in the actual world. But could it possibly happen?

We can imagine a case of testimony: someone walks up and tells us "You need to save this child. Trust me". If that person has a child-detector that tells her whether there are other children drowning in the lake, then I could reasonably infer that there aren't one hundred children drowning in the lake. Similarly, if she is better than me at metaphysics, I would think that if she's so sure I should save this child, then she must have concluded that modal realism is false, and so I do gain justification for the claim that modal realism is false. But this isn't genuine Moralism. Ultimately, her knowledge that I should save this child is inferential on her knowledge of the fact that there aren't those other drowning children, and that modal realism is false. It's just that she is empirically and metaphysically better informed than I am, and I indirectly inherit that information through her act of testimony. She could have just directly said that there aren't those drowning children, and that modal realism is false. And even if this was genuine Moralism, it's irrelevant in the context at hand. The question is whether the fact that a metaphysical theory has a counterintuitive ethical consequence is a reason against it. Inasmuch as the consequence is counterintuitive, it's because it seems to us to be counterintuitive, not because someone told us so.

Maybe, stretching our imagination a little bit, we can imagine genuine cases of Moralism. Maybe God directly knows all the facts, including the first-order ethical ones. So, God could "infer" a metaphysical theory from this ethical claim. But since he already knows the true metaphysics, it's unclear what the point of that would be. Or we can imagine that we had a "moral crystal ball": a device that spits out true ethical claims. We question the ball, and it says that we ought to save the child. Then plausibly, we get direct justification for this claim, and we can go on to infer that there aren't other drowning children, and that modal realism is false. But I can't really imagine how this moral crystal ball works. The only way I can make sense of it is if it somehow knows the descriptive facts and infers the first-order ethical facts from them. But then, this is just like the case of testimony above. And at any rate, no crystal ball is telling us ethical facts that could allow us to rule out modal realism.

In conclusion, there may or may not be possible cases of successful Moralist arguments. But all the Moralist arguments that we actually make will fail. Ethical justification is direct only insofar as it pertains to moral laws or particular moral conditionals; ethical justification in first-order ethical claims is always inferential.

And so, modal realism is saved from the ethical objection. And in general, we cannot reject a metaphysical theory because it has a strange first-order ethical conclusion. But there are other kinds of inferences from ethics to metaphysics. Let's now see if Anti-Moralism is true of them as well or not.

6 Metaphysical Notions that are Ethical in Nature

Some metaphysical notions are essentially intertwined with ethical notions. For example, notions like personal identity, free will, agent causation, seem essentially tied to facts about blame, responsibility, punishment, prudential rationality, and so on. One might naturally think that in *those* cases, at least, the inference from ethical premises to metaphysical conclusions will be legitimate. Barber (2020, *p*. 483), for example, writes that in cases involving such notions, metaphysics is "uncontroversially susceptible to moral refutation".[16]

I will argue that that's not the case, by focusing on personal identity. The question of personal identity is the question of what the conditions are under which a person persists through time. On some views, personal identity requires a certain kind of psychological connection between the person at different times, involving for example memory. On other views, it requires a physical connection, involving the continuous existence of the body or of the brain.[17]

[16] In general, Barber claims (although his terminology is different from mine) that when a metaphysical notion is naturalistic (that is, continuous with science), Anti-Moralism is true; but when it's anti-naturalistic, then Moralism is true; and normative notions, being anti-naturalistic, license good ethics-to-metaphysics inferences. What I will argue in the rest of this section is that an additional distinction needs to be made among the normative notions, using the notion of metaphysical dependence. Depending on which facts metaphysically depend on which facts, it will turn out that while some normative notions are indeed anti-naturalistic in his sense (what I will call the "strongly normative notions"), others are actually naturalistic (what I will call the "weakly normative notions").
[17] Defenders of the psychological theory include Lewis (1976), Parfit (1984). Its contemporary form descends from Locke. Proponents of physical views (in one form or another) include Thomson (1997), Olson (1997).

The notion of personal identity is of special normative importance, in that it appears in a network of normative claims. For example, people ought to be blamed for things *they* did at earlier times; or we may have special reasons to care about what will happen to *us*, as opposed to other people. As John Locke said, "Personal identity is a forensic notion"; and many have followed him. If so, one might think that we can draw conclusions about personal identity from ethical premises, and, if need be, refute a theory of personal identity on ethical grounds.

Can we? Let's try it out. Let's make up a theory of personal identity with wild ethical consequences. For example, consider the *Waking theory of personal identity*. It says that a person only survives while they're awake. When they go to sleep, they die, and each morning there is, strictly speaking, a new person (inhabiting the same body, and with similar memories, et cetera).

Regardless of its merits, this theory exemplifies the main topic of this paper. This is because the Waking theory entails many revisionary ethical claims. As we saw before, we typically think that people ought to be blamed (and punished) for what they do. Take Joe, who is in jail for brutally murdering several people back in 2018. It is pretty obvious that he ought to be in jail. At least, that's what we think. But if the Waking theory is right, Joe didn't commit any crime: Joe came into existence this morning when he woke up. Someone else, years ago, committed those crimes. Joe would be actually innocent, and blameless, and arguably should be let out of jail. So, here comes the Moralist argument:

P1 Joe ought to be blamed and punished;
P2 If the Waking theory is right, Joe ought not to be blamed and punished; therefore,
C The Waking theory is false.

Let's call this the Waking argument. Is this Moralist argument any good? The thought behind this section is that this argument, unlike the Lake argument, is a good one, because personal identity has an ethical nature.

I will argue that the Waking argument is not a good argument, and personal identity no counterexample to Anti-Moralism. Whether this is true hinges on some metaphysical facts involving personal identity, that can be illustrated using a different example where Moralism is actually true.

Consider the notion of "personal space". It's a metaphysical notion – a region of space surrounding a person – which is clearly ethical in nature. For simplicity's sake, let's suppose this ethical nature is fully expressed by the following condition: space region X is someone's personal space iff it's *prima facie* wrong to enter X without the person's consent.

In this context, it seems that we *can* legitimately infer metaphysical facts from ethical premises. For example, we have the intuition that we shouldn't get within one inch of someone without their consent. From this, we can infer that their personal space spans at least one inch from their body. And if a certain metaphysical theory T had the consequence that someone's personal space spans only half an inch, then our intuition is definitely some reason against this theory, maybe even conclusive reason. So in this case, Moralism seems right. Why is it, and how does it differ from the diagrams shown before?

The criterion of transmission failure defended in section 3 said that if our justification for believing a proposition P is fully inferential on its conditions, then we cannot infer any of its conditions from P. This criterion doesn't apply in the case of personal space, because "X is someone's personal space" is not a condition of "It's prima facie wrong to enter X". Why is that? The conditions of something, in the sense relevant for the criterion, are the things that are metaphysically prior to it. In the Lake example, the descriptive facts of the case (that a child is drowning, that we can make a difference, etc.) are metaphysically prior to the mixed ethical fact that we should save this child (that is, we should save the child *because* a child is drowning, etc.); and that's a requirement for it to be the case that we cannot infer the descriptive facts from the mixed ethical fact. In the case of personal space, this requirement doesn't hold: the fact that X is someone's personal space is not metaphysically prior to the fact that it's wrong to enter X. In fact, the opposite is true: X is someone's personal space *because* it's wrong to enter X. That's the sense in which personal space is a metaphysical notion that is ethical in nature: the facts involving that notion are made true by the ethical facts.[18]

Ultimately, then, the key difference between personal space and the cases discussed earlier is a difference in metaphysical priority, which then translates into a

[18] When I say "metaphysically prior", I mean whichever relation one think plays the role of giving structure to reality. In the ethical cases in question, one could mean reduction or grounding, and interpret the various relevant claims accordingly.

difference in justificatory priority. We are justified in believing claims about someone's personal space inferentially, from ethical claims about whether we may or may not enter some region of space.[19] This is why in this case, we can infer a metaphysical fact (and reject a revisionary metaphysical theory) on the basis of ethical facts.

OK, so Moralism is true of personal space, and we now know why. But that's not a very interesting notion; what about the original example, personal identity?

The discussion above shows that the crucial question is whether the facts about personal identity are metaphysically prior or posterior to the related ethical facts about punishment. In the example of Joe above, the identity fact is "Joe is the person who committed the crime", while the ethical fact is "Joe should be punished". If the punishing facts were prior to the identity facts, then personal identity would be just like personal space. The identity facts wouldn't be conditions of the punishing facts, and so the Waking argument wouldn't be a case of transmission failure. Moralism would be true of arguments involving personal identity.

Suppose, instead, that the identity facts are metaphysically prior to the punishing facts: that is, suppose that we should punish Joe *because* he is the person who committed the crime. Then the case of personal identity would not be like personal space, but would rather be like the case of modal realism. The identity facts would be conditions of the punishing facts, and so the Waking argument would be a bad argument, and Anti-Moralism would be true of arguments involving personal identity.

We can illustrate the difference between these two views by considering an "easy case" of Anti-Moralism involving facts about punishment. Going back to Joe, there are many empirical possibilities in which he is, in fact, innocent. Suppose the person who committed the crime (the original Joe), and was put in jail since, dissolved into thin air overnight, and, by chance, a perfect duplicate (the person we now call Joe) materialized out of nowhere into his cell. Then, we wouldn't have to punish Joe. So, here's the "materialized overnight" argument:

P1 Joe ought to be blamed and punished;

[19] Of course, our belief in this ethical fact is itself justified from other descriptive facts. The fact that we cannot get within one inch of someone might be inferred from, for example, facts about the consequences of such an action.

P2 If Joe materialized overnight out of nowhere, Joe ought not to be punished; therefore,

C Joe didn't materialize overnight out of nowhere.

Clearly, we can't be justified in the conclusion on the basis of this argument. And the diagnosis for this is the now familiar one. We're justified in thinking that Joe should be punished on the basis of its conditions, one of which is that he didn't materialize overnight; and so, we can't infer the latter from the former. If one thought, in addition, that the identity facts are also conditions of the punishing facts, that is, that Joe should be punished because he is the same person as the person who committed the crime, then we cannot also infer facts about personal identity from facts about punishment; the Waking argument would be just as bad as this argument, and for the same reason. Those, instead, who think that punishing facts are prior to the identity facts, will think that the Waking argument is different in structure from this argument.

The question, then, is whether the identity facts are metaphysically prior to the punishing facts, or vice-versa. That is, should Joe be punished *because* he is the same person as the murderer? Or is he the same person as the murderer *because* he should be punished? I will argue that Joe should be punished because he is the same person as the murderer; and so, that identity facts are conditions of the punishing facts; and so, that the Waking argument is not a good argument, and that Anti-Moralism is true of personal identity. I have three reasons in favor of this.

The first reason is intuition. The natural thing to say is that Joe should be punished because he is the same person as the murderer. If we ask why he should be punished, that would be the obvious thing to say. The alternative is to say that Joe is the same person as the murderer because we should punish him; but that seems strange, and appears to get things in the wrong order.

The second reason is related to the first, and has to do with providing good explanations for the punishing facts. Why should we punish Joe? On my view, we should punish Joe because he is the murderer. That's a simple, natural, and intelligible explanation. On the alternative view, we cannot say that, because the identity fact is actually explained by the punishing fact. Rather, on the alternative view we could only say something like "we should punish Joe because he is a continuous organism with the murderer" or "we should punish Joe because his mental states are relevantly similar and caused by the mental states of the murderer",

depending on which criterion of personal identity one likes. But such an explanation seems to me less intelligible than the identity explanation. I do see why being the same person as the murderer explains why we should punish Joe; I don't see how being a continuous organism with the murderer (say) explains why we should punish Joe. If being a continuous organism is relevant to punishing, it's only because being a continuous organism with the murderer is sufficient for being the same person as the murderer, which in turn explains why we should punish Joe; but then, we need the personal identity facts in the middle; and so we go back to the view that identity explains punishment.

The third and final reason is that if the identity facts were grounded in the ethical facts, then in the absence of ethical facts the notion of personal identity would be empty. Yet, that doesn't seem the case. Even in the moments in which I'm tempted by moral anti-realism, I still think that there are facts about personal identity. I can first personally wonder whether I can survive Parfit-style teletransportation even if I thought there were no true ethical facts. In the case of Joe, I can similarly believe that there are no moral facts at all, so that he shouldn't be punished, yet still wonder whether he is in fact the murderer. That would be nonsensical if what makes Joe the same person as the murderer were ethical facts involving punishment or the like. Crucially, the opposite happens with the notion of personal space discussed earlier. Inasmuch as I feel tempted by moral anti-realism, I do not feel like there are any facts about personal space left wondering about; the notion of personal space becomes empty if there is no morality. And that's because facts about personal space are grounded in ethical facts.

These are my reasons for thinking that the identity facts are prior to the punishing facts. If this is right, then our justification for believing that Joe should be punished is inferential on our justification for believing that he is the murderer; hence, the Waking theory cannot be rejected because it has the strange ethical consequence that we shouldn't punish Joe. That might seem strange: isn't it obvious that we should punish Joe? Yes, it is obvious; but it's obvious only because it's antecedently obvious that the Waking theory is false and that we survive sleep.

All is well. But there's a puzzling sensation that remains. In what sense is personal identity a normative notion, if it doesn't license good inferences from ethics to metaphysics? Here we need to draw a distinction between two different kinds of normative notions. A metaphysical notion is *strongly normative* when the facts

about it hold in virtue of ethical facts. Personal space is an example of a strongly normative notion, since X is someone's personal space because it's wrong to enter X. Personal identity, I've argued, is not a strongly normative notion. However, it's still a normative notion, just in a weaker sense. A *weakly normative* notion is such that, while the facts about it don't hold in virtue of ethical facts, it still plays a central role in our first-order normative theorizing. Personal identity does play such a central role, by entailing punishing facts (and many other ethical facts), so personal identity is a weakly normative notion. It's still a normative notion, unlike other notions which aren't normative even in a weak sense (like the notion of abstract object, which doesn't play a central role in our normative theorizing). Moralism is true of arguments involving strongly normative notions, but it's false of arguments involving weakly normative notions.

In conclusion, in spite of appearances, personal identity is no counterexample to Anti-Moralism. Of course, you might disagree. You might think that personal identity is a strongly normative notion, like personal space. This seems wrong to me, for the reasons given above. But at any rate, I've pointed out the criterion by which we can ascertain whether we can infer a metaphysical conclusion from an ethical premise: the facts about the metaphysical notion must hold in virtue of the ethical facts. This is true of some notions, like personal space; I've argued that it's not true in the case of personal identity, which was probably the most interesting candidate for Moralism.[20] On the whole, then, Anti-Moralism is still true.

7 The Metaphysics of Pure Ethics

The final category of ethics to metaphysics inference I want to consider concerns inferences in which the metaphysics is relevant to the purely ethical facts in question. This happens, for example, when it's the metaphysics of moral facts themselves. Consider the notion of "moral property", such as the goodness or badness of an action or of a state of affairs. It's a metaphysical notion, but it's directly about ethics.

[20] In the case of free will, one would need to figure out whether the facts involving free will, such as "Joe's action X was free", are prior or posterior to the related ethical facts, such as "Joe is responsible for action X". Here too it seems (although I won't argue for it) that Joe is responsible *because* his action was free (and not viceversa), and so that free will is a weakly normative notion.

When a notion such as this one is involved, we can get a Moralist argument that might be good. Suppose we thought that the existence of God is required for there to be moral properties and moral facts. Going back to our lake, an atheist could give the following argument:

P1 God doesn't exist;
P2 If God doesn't exist, I don't have to save the child; therefore,
C I don't have to save the child.

And here is the flipped, Moralist argument:

P1 I have to save the child;
P2 If God doesn't exist, I don't have to save the child; therefore,
C God exists.

Many have been tempted by an argument for God's existence along these lines.[21] Is it any good?

Our justification for "I have to save the child" is inferential on the descriptive facts of the case, as I've argued in the previous sections. There is, however, a fact in the vicinity such that the justification we have for it is arguably not inferential: the relevant pure normative fact. For the generalist, it's the fact that we ought to follow this or that moral principle; for the particularist, it's the fact that *if* things are such-and-such (a child is drowning, etc.), then I have to save him. If P2 in the argument is true, then God's existence is a condition for this purely normative fact to be true. If our justification to believe this normative fact is wholly inferential on its conditions, then, we couldn't possibly infer from it that God exists. If, instead, we have direct justification to believe it, then we *can* infer from this fact that God exists. The Moralist argument, then, is really the following:

P1 If a child is drowning (etc.), I have to save the child;
P2 If God doesn't exist, P1 is false; therefore,
C God exists.

The question is then whether P1 is justified inferentially, or directly. The followers of moral perception, or those who say that we intuit general moral principles directly, will say that it is justified directly, and so that this argument is successful. Others might say that we're only justified in thinking that it *appears* that we have moral

[21] See Evans & Baggett (2014).

obligations, unless we're justified in the conditions for that claim, such as the existence of God and of moral properties.

When a metaphysical claim is directly relevant to pure ethics (as opposed to the mixed facts), then Moralism may be the right answer, conditional on the directedness of our justification in pure ethical claims. Even if that were true, that wouldn't do much by itself to overturn a metaphysical theory. The real meat of the argument lies in P2: the claim that God is required for morality. That is an extremely controversial premise, and so the possibility that this ethical premise leads us to this metaphysical conclusion is remote.[22] This is in contrast with the case of modal realism (and other arguments similar to that), where there the premise that links the metaphysical theory to a revisionary ethical consequence is more plausible.

Conclusion

Let's sum up. We've seen three different kinds of arguments from ethical premises to metaphysical conclusions. In the first kind, we infer a metaphysical fact from a mixed ethical fact. I've argued that these arguments are epistemically circular and fail to transmit justification, and so Anti-Moralism is true of them. In the second kind, we use an ethical fact to infer metaphysical facts involving metaphysical notions that are ethical in nature. I've argued that these arguments are good, but they require a strong conception of "metaphysical notions ethical in nature", and the most interesting and plausible candidate (personal identity) doesn't satisfy such a conception. In the third and last kind, we infer a metaphysical fact from a purely normative fact. In this case, Moralism might be true, conditional on a certain view about our justification we have for believing the purely ethical facts; but in the case I've considered (the moral argument for God) it relies on a very controversial second premise. All in all, then, when it's ethics vs. metaphysics, metaphysics stands.

[22] For another example of this kind of argument, suppose someone thought that physicalism (the claim that everything is, or supervenes on, or is grounded on, the physical) is incompatible with the existence of moral facts. Then one could gain reason to believe that physicalism is false on the basis of P1. And again, the premise that is hard to defend is precisely the linking premise that physicalism is incompatible with the existence of moral facts. For a last example where similar considerations apply, one can consider the argument from P1 to moral objectivism.

References

Adams, R. (1974), "Theories of Actuality", *Noûs*, 8(3): 211–231.

Baggett, D., and Walls, J., (2019), *The Moral Argument: A History*, Oxford: Oxford University Press.

Barber, A. (2020), "Is Metaphysics Immune to Moral Refutation?", *Acta Analytica* 35 (4):469-492.

Davies, M. (1998), "Externalism, Architecturalism, and Epistemic Warrant", in *Knowing Our Own Minds*, Crispin Wright, Barry C. Smith, and Cynthia Macdonald (eds.), Oxford: Clarendon Press, 321–362.

Evans, C. S. and Baggett, D. (2014) "Moral Arguments for the Existence of God", *The Stanford Encyclopedia of Philosophy (Winter 2022 Edition),* Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives /win2022/entries/moral-arguments-god/>.

Hawthorne, J. (2004), *Knowledge and Lotteries*. Oxford: Clarendon Press.

Heller, M. (2003), "The immorality of modal realism, or: How I learned to stop worrying and let the children drown", *Philosophical Studies* 114 (1-2):1 - 22.

Johnston, M. (2016), "The Personite Problem: Should Practical Reason Be Tabled?", *Noûs* 50 (4):617-644.

Lewis, D. (1976), "Survival and Identity", in A. Rorty (ed.), *Identities of Persons*, Berkeley, CA: University of California Press; reprinted in his *Philosophical Papers vol. I*, New York: Oxford University Press, 1983.

——— (1986), *On the Plurality of Worlds*, Oxford: Blackwell Publishers.

Parfit, D. (1984), *Reasons and Persons*. Oxford: Clarendon Press.

Pryor, J. (2004), "What's Wrong with Moore's Argument?", *Philosophical Issues*, 14(1): 349–378.

McGrath, S. (2004), "Moral knowledge by perception", *Philosophical Perspectives*, 18(1), 209–228.

Moretti, L. and Piazza T. (2023), "Transmission of Justification and Warrant", *The Stanford Encyclopedia of Philosophy* (Summer 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/ sum2023/entries/transmission-justification-warrant/>.

Olson, E.T. (1997), *The Human Animal: Personal Identity Without Psychology*, New York: Oxford University Press.

Ridge, M. and McKeever S. (2016), "Moral Particularism and Moral Generalism", *The Stanford Encyclopedia of Philosophy (Summer 2023 Edition)*, Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/sum2023/entries/moral-particularism-generalism/>.

Silva, P. & Oliveira, L. R. G. (forthcoming), "Propositional Justification and Doxastic Justification", in Maria Lasonen-Aarnio & Clayton M. Littlejohn (eds.), *Routledge Handbook of the Philosophy of Evidence*. Routledge.

Stoljar, D. (2001), "Physicalism", *The Stanford Encyclopedia of Philosophy (Summer 2023 Edition)*, Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/sum2023/entries/physicalism/>.

Thomson, J. J. (1997), "People and their Bodies", in *Reading Parfit*, J. Dancy (ed.), Oxford: Blackwell.

Tucker, C. (2010), "Transmission and Transmission Failure in Epistemology", *Internet Encyclopedia of Philosophy* 1.

Vogel, J. (1990), "Are There Counterexamples to the Closure Principle?" in *Doubting: Contemporary Perspectives on Skepticism*, M. Roth and G. Ross (eds.), Dordrecht: Kluwer Academic Publishers.

Werner, PJ. (2020), "Moral perception", *Philosophy Compass*.

White, R. (2006), "Problems for Dogmatism", *Philosophical Studies*, 131(3): 525–557.

Wright, Crispin, 1985, "Facts and Certainty", *Proceedings of the British Academy*, 71: 429–472.