

The Worst Possible Misery for Anyone: A Close Reading of Sam Harris' Science of Morality

Mark Douglas Warren

Sam Harris thinks science can answer moral questions.

There are plausible ways to understand this: You might think he's advocating that our judgments about what to do should be informed by an understanding of the world as it's revealed by science. And that's just obviously true, the kind of thing that's standard operating procedure in most people's ethical reasoning. If you want to decide whether abortions are ethically permissible, you might want to know about the science of brain development in fetuses. If you're trying to decide if the death penalty is just, you should at least hear what sociologists have to say about racial disparities in capital punishment. To figure out what to do, we need to know how things are with the world, and science is very good at filling in those blanks.

But Harris is saying something more than that. Science won't just *help* us find moral answers; science can give those answers. Harris wrote a best-selling book making the argument—*The Moral Landscape: How Science Can Determine Human Value*. He intends his readers to take that subtitle seriously. Indeed, he thinks morality should itself “be considered an undeveloped branch of science.” (Harris 2012, 4) Science has an impressive track record of answering some of our most vexing questions. If Harris is right, we will one day add, “How should I live?” to that list.

Let's see if he's right.

Living Better Through Chemistry?

To many, Harris' position will seem like a gross exaggeration of science's proper role. This objection can be traced back to the 17th century philosopher David Hume's famous is-ought problem. (Hume [1739] 2007) He argued that no amount of information about how the world *is* can, on its own, lead us by force of logic to a conclusion about how things *ought* to be. Science is in the “is” business; morality is in the “ought” business. It follows that science can only deliver oughts in an advisory capacity.

Should I take this pill? If I want to know what's in it before I decide, someone trained in the science of chemistry can tell me if, say, it's vitamin D. Or if it's MDMA. And someone trained in the medical sciences can tell me, “If you want healthy bones, you should take the vitamin D.” Or: “If you want to dance for three hours, you should take the Ecstasy.” The advice I can get from scientists like these are contingent on my aims. Science can tell me, if I want a particular outcome, here's the way I should proceed.

But my question might run deeper. Perhaps I already know what's in the pill, but I don't know what outcome I *should* want. Should I strive to maintain robust health? Should I throw caution to the wind and seize the day? Is it better to prevent osteoporosis, or to get really high? Conventional wisdom would tell us that such questions are outside of the purview of science.

Harris disagrees with Hume, and the conventional wisdom that followed him:

I am arguing that science can, in principle, help us understand what we should do and should want—and, therefore, what other people should do and should want in order to live the best lives

possible. My claim is that there are right and wrong answers to moral questions, just as there are right and wrong answers to questions of physics, and such answers may one day fall within reach of the maturing sciences of mind. (Harris, 28)

Before I move on to assess this claim, we should pause to appreciate just how revolutionary it is. If it succeeds in its arguments, this book is an incredible accomplishment. It would overturn the dominant ideology about the relationship between science and values. It would bridge the divide between “is” and “ought”.

From Science to Morality in Four Easy Steps

How do you get from the raw material of the scientific inquiry to moral insight? Here are the steps Harris takes:

1. The well-being of conscious creatures is completely dependent on environmental and neurological states.
2. These states are scientifically tractable.
3. Moral truths are entirely reducible to truths about the well-being of conscious creatures.
4. Therefore, science is uniquely positioned to answer moral questions.

If you’re paying attention, that third premise should jump out at you. It’s a passable summary of what philosophers would call a *utilitarian* worldview. Utilitarianism is the moral philosophy that the rightness or wrongness of an action boils down to whether and how much happiness or suffering the act causes. Every time I teach an ethics course, I devote a few days of class to explaining the utilitarian philosophy, drawing out its implications for how we should behave, comparing it to other moral philosophies, and rehearsing the challenges it has faced from skeptical philosophers over the past couple of centuries.

So, you’d think this particular premise would belong to the philosophers, and not the scientists. *The Moral Landscape: How Science Can Determine Human Values (if Utilitarianism Turns Out to Be True)* doesn’t have quite the same ring to it, but perhaps it would be a more accurate title.

This objection is worth considering carefully. Harris presents us an argument that science can determine our values. He’s promised a bridge that can take us from the “is” of science to the “ought” of morality. But it looks like he’s packed all the load-bearing oughts of his argument in that third premise, and *that* premise looks more like philosophy than science. So, a skeptic might at this point ask him: “Is this moral principle itself supposed to be a scientific truth?”

Scientists Can Help Themselves to Self-Evident Premises

But Harris isn’t bothered by this question. He argues that philosophy and science shouldn’t be thought of as entirely distinct enterprises—“science is”, he claims, “often a matter of philosophy in practice.” (Harris, 180) Because of this, he doesn’t think it’s such a great sin for a branch of science to smuggle in philosophical assumptions. In fact, he argues, it’s unavoidable: “It is essential to see that the demand for radical justification leveled by the moral skeptic could not be met by any branch of science.” (Harris, 37)

Consider for example the study of medicine. Doctors work on the assumption that their job is to promote their patients’ health, broadly construed. But if a doctor were asked how they could possibly justify this assumption, Harris would, I think, recommend a hearty rolling of the eyes as

an adequate response. “Science cannot tell us why, scientifically, we should value health. But once we admit that health is the proper concern of medicine, we can then study and promote it through science.” (Harris, 37)

The argument, then, is that just because his argument relies on a non-scientific premise, that alone doesn’t undermine the scientific bona fides of his conception of morality. *All* scientific enterprises rely at some point on non-scientific premises. All scientists, for example, assume that regularities we observe in the past will hold in the future, that logic is universal, that truth is worth pursuing. But we don’t have to use science itself to prove these assumptions true in order to engage productively in the scientific enterprise. And that’s because these principles are just self-evident. So yes, Harris’ science of morality relies on a non-scientific premise, but that’s not enough to disqualify it as a science. It’s in good company.

Now, dear reader, in the spirit of charity, I’m going to ask you to ignore a lot. Ignore any worries that Harris has conflated medical *science*—the aim of which is to understand how the human body works, and how disease and other maladies operate—with the *practice* of medicine—the aim of which is to use findings from medical science to promote human health. Ignore any concern you might have that whether his argument counts as science or not, it still hasn’t met Hume’s challenge. Ignore any sneaking suspicion that, by muddying the distinction between philosophy and science, Harris has given himself just enough cover to do good old fashioned moral reasoning and call it science.

We’ll just grant all these points to Harris. I’m going to ask you to focus instead on a problem that lies at the heart of his argument: Harris’ utilitarian premise is *not* self-evident in the same way that the other assumptions we’ve talked about are. Any scientist worth their salt thinks that truth is worth pursuing. Any doctor worth visiting will agree that health is worth promoting. But moral philosophers disagree vehemently about Harris’ moral premise—utilitarianism is, in fact, a minority position among academic philosophers.

This objection is not dispositive. The majority of philosophers might just be wrong. They might miss a point that should be self-evident. But to deal with this concern, Harris will need to show us *why* they should be ignored. And he’s got an argument to that effect. To make his point, he’ll have to bring us all to hell.

The Worst Possible Misery for Everyone

Imagine a world of perfect suffering. The universe God would make, if His benevolence soured into hatred: Every conscious creature in this world suffers as much as its neurology will allow, every moment of the day, for all of time. Imagine hell on earth.

Such a world would be...well, pretty bad, right? This is self-evident. It would in fact be as bad a situation as we can imagine. Once we admit this, Harris argues, we’ve accepted the first link in a chain of logic that leads us inexorably to his utilitarian premise:

Once we conceive of “the worst possible misery for everyone,” then we can talk about taking incremental steps toward this abyss: What could it mean for life on earth to get worse for all human beings simultaneously? It seems uncontroversial to say that a change that leaves everyone worse off, by any rational standard, can be reasonably called “bad,” if this word is to have any meaning at all. (Harris, 39)

The idea is that, if you agree that the worst possible misery for everyone is bad, and that any improvement in the well-being of everyone would make things better, you’ve admitted that

questions of goodness and badness must ultimately be ones about how things are going for sentient creatures.

I am saying that a universe in which all conscious beings suffer the worst possible misery is worse than a universe in which they experience well-being. This is all we need to talk about “moral truth” in the context of science. (Harris, 39)

Let’s recap. Harris thinks that science can answer moral questions. His argument for that position rests on the non-scientific assumption that morality is basically about promoting the well-being of conscious creatures. This isn’t a problem, he thinks, because by reflecting on how bad an imagined hell world would be, we can see the obvious connection between what we value and the experiences of conscious beings. Once we admit this, we must also admit that morality is fundamentally concerned with promoting well-being, with moving us as far away from such a world as possible.

And all of this just as self-evident as any other non-philosophical assumption made by the sciences. So, inasmuch as the well-being of conscious creatures correlates to their brain states, and their brain states can be studied by science, we’ve laid the foundation for a science of morality.

Adding a Bit of Misery

I don’t think this is the right way of framing the issue, for reasons I’ll explain in a bit.

But for now, it’s worth noting that this argument doesn’t even work on its own terms. To see why, imagine another world, an afterlife. Its denizens are those men and women who spent their time here on Earth in the vilest ways imaginable—they are the tyrants, the predators, the social media tycoons of this world. But surprisingly, their hereafter isn’t so bad, all things considered. It’s something like life for a middle-class American. Not so great, but not too bad, either.

Now imagine you had the opportunity to move that world one iota nearer Harris’ miserable extreme. You can push a button and by doing so introduce one small uncompensated misery to every villain in this afterlife. A mouse bites Hitler’s toe one day. Pol Pot catches a cold. Stalin’s wife keeps cutting out.

Would you push the button?

If Harris is right about the self-evidence of utilitarianism, the question should seem ridiculous. It’s like asking a doctor if they should try to help their patient, or a physicist if we should believe in an external world. It should just be *obvious* to us that it would be wrong to intentionally cause harm for its own sake.

But even if you wouldn’t push the button, you might be tempted, right? In one sense, doing so would unambiguously make that world a worse place—worse for each person in it. But it’s not clear that doing this would make the world a *morally* worse place. It’s just wrong, you might think, that all these wretched people get to go on living their afterlives as if they hadn’t been monsters. It’s not fair! And you needn’t fantasize about eternal punishment in a fiery pit to think that perhaps the damned at least deserve some fair share of the misery they caused the rest of us.

This is a distinction that Harris misses: we can admit that things can go better or worse for people without thereby committing ourselves to a judgment about whether such eventualities are *morally* good or bad. His arguments here conflate badness—as an assessment of how things are for someone—with *moral* badness. We can concede that something is bad in the first sense (e.g.,

Pol Pot is having a hard time with his sinus congestion) without drawing any direct conclusions about its moral status.

That's because considerations of morality outstrip considerations of how things are going for people. Contra Harris, there are other things we care about, other things worth valuing. If it's unclear to you whether you should push the button, your uncertainty reflects this. Justice is another thing worth valuing, and a concern for justice can't always be neatly translated into a concern for well-being.

And justice isn't the only competing moral value. Over the millennia, philosophers have enumerated many others: autonomy, virtue, equality, fraternity. Harris briefly acknowledges some of these alternate conceptions of morality throughout the book. In each case, his counterargument is the same: These considerations might *look* like they speak against a monolithic utilitarian account of moral value, but if we're honest with ourselves, we will admit that ultimately justice, autonomy, etc., can all ultimately be reduced to concerns about well-being.

For example, in his argument against justice as an independent source of moral value, he admits that justice can be very useful, inasmuch as just societies tend to be happy societies. But he asks us:

How would we feel if, after structuring [an ideally just society], we were told by an omniscient being that we had made a few choices that, though eminently fair, would lead to the unnecessary misery of millions, while parameters that were ever-so-slightly less fair would entail no such suffering? Could we be indifferent to this information? (Harris, 79)

The implication is that in those situations where considerations of justice are in competition with considerations of well-being, it should just be obvious to the reader that well-being should win out.

It's not obvious to me, at least. His thought experiment can be neatly turned on its head: Imagine a society where everyone is fairly well off, but we've found a way to increase net happiness a small amount. Unfortunately, doing so will exacerbate existing injustices: inequalities will be heightened, bad actors will be rewarded, legitimate grievances will go unanswered. Is it just self-evident that we shouldn't care, that justice must be sacrificed at the altar of beneficence?

But Harris never considers such a situation. He's got a utilitarian hammer, and every alternate moral value looks like a nail. I will leave it as an exercise for the reader to decide whether Harris' other counterarguments in this vein are convincing. I have bigger fish to fry.

The Worst Possible Mechanic

It is not self-evident that every incremental move toward more suffering makes for a world that is morally worse. Not necessarily, anyway. And by the same token, it's not self-evident that every move towards greater well-being makes for a world that is morally better. But I claimed above that this doesn't get to the heart of the matter. Harris' argument suffers from a deeper flaw.

I can illustrate this point if you'll grant him all the points he's made so far. We've been so charitable already; let us go a bit further. Grant Harris that not only is suffering bad, but it is also morally bad, and necessarily so. Grant him that an increase in suffering automatically makes a situation morally worse, and that an increase in well-being makes it morally better. Grant him

that all this argumentation can plausibly be thought of as a seamless part of a properly scientific process.

Let our charity reach a fevered pitch. Grant him everything he's asked that you admit to yourself, even if, like me, you're skeptical you should.

Because once you grant him all this, then you can see the real flaw in his argument: even if we accept that suffering is morally bad and flourishing is morally good, we still haven't shown that suffering is *all* that's bad, and we definitely haven't proven that all there is to badness is misery, or that all there is to goodness is well-being. A world of maximal suffering might be the worst moral situation possible, but it could still be the case that moral questions cannot be reduced to questions of suffering and well-being.

To see this, imagine you bring your car into the shop one day. The wheels need alignment. Your oil needs changing. You can't get the radio to play. The mechanic confidently opens the hood to the car, sprays half a can of Rust-Oleum over the engine, closes the hood, and presents you with a bill.

Of course, you're outraged and refuse to pay up. Rust-Oleum isn't going to fix your radio! He didn't even change your oil! But in the sonorous tones of a veteran podcaster, he tries to calm your worries. He explains: "Imagine a car whose parts have rusted through entirely. Every square inch of metal on this vehicle has been eaten to nothing but brittle and hole-pocked rust. It's perfectly oxidized. Now, if you have any opinion worth listening to about automotive issues, it will just be clear to you that this car is not going to run—that indeed such a situation represents a unique nadir of motor functionality. Any addition of rust to a healthy car will make it worse off; any removal of rust will make it better. It follows, then, that questions of a car's automotive health admit of right and wrong answers, and these questions are reducible to questions about the oxidation of its parts."

Now, you might have doubts about whether a completely rusted out car is the worst possible car, but you should ignore those for the moment. Because that just distracts from the more urgent issue at hand: even if he's right, it still doesn't follow that fighting rust is the only way to promote automotive health. Just because rust is bad doesn't mean it's the *only* bad thing that can happen to a car.

Science Alone Cannot Determine Human Values

Readers who bought Harris' book in the hopes of understanding how science can answer moral questions should be equally disappointed. Even if we admit that the worst possible misery for everyone is the worst moral situation possible, even if we agree with Harris that anything that brings us closer to such a situation must be bad—and morally bad, at that—we still haven't shown that *all there is* to morality is the promotion of well-being and the abatement of suffering.

The failure to show that his utilitarian premise is self-evident is disastrous for Harris' project. He needs to secure this point if he wants to plausibly put morality alongside the other sciences. Without it, he's stuck with the rest of us as we navigate the difficult question of how we should live. There are inescapably philosophical issues with which we must grapple if we hope to do this successfully. We must weigh sometimes competing values—justice, beneficence, autonomy, virtue—to make up our minds about which are the most appropriate to act on in a given situation.

And yes, the best moral deliberators are using science where it's relevant to their principles. But they're not *doing science* when they decide which principles are important. A chemist can't tell

you how to live your life; for that, you'll have to go to a priest or a philosopher. Science is in the business of telling us how the world works, but we need something else—religion or philosophy or poetry—to tell us how the world should be, and therefore how we ought to behave.

Harris, Sam. *The Moral Landscape: How Science Can Determine Human Values*. London: Black Swan, 2012.

Hume, David. *A Treatise of Human Nature*, ed. D. F. Norton and M. J. Norton. 1739. Oxford: Oxford University Press, 2007.