

## REVIEW OF *THE MORAL LANDSCAPE*

Two quotes which I think sum up Harris's argument:

It seems clear that what we are really asking when we wonder whether a certain state of pleasure is "good," is whether it is conducive to, or obstructive of, some deeper form of well-being.<sup>1</sup>

...human well-being entirely depends on events in the world and on states of the human brain. Consequently, there must be scientific truths to be known about it.<sup>2</sup>

Harris is (I think) writing primarily for a specific subculture of educated, liberal atheists who regard science as the most credible source of guidance in life. Some people in this subculture are very hesitant to condemn any aspect of other cultures as evil. One reason they give for this hesitancy is that the whole notion of objective moral truth is incoherent; in other words, they're moral relativists. Harris wants to convince those people that objective moral truth exists, so that they will help resist a variety of practices that cause unnecessary suffering around the world.

He categorizes his view as a form of [consequentialist moral realism](#) (note: "realism" in this context has nothing to do with cynicism or pragmatism):

...this commits me to some form of moral realism (viz. moral claims can really be true or false) and some form of consequentialism (viz. the rightness of an act depends on how it impacts the well-being of conscious creatures).<sup>3</sup>

I support that conclusion but I think there's an important gap in Harris's argument for it. The gap is identified by G. E. Moore's [Open-question argument](#), which Wikipedia summarizes as follows:

Premise 1: If  $X$  is good by definition, then the question "Is it true that  $X$  is good?" is meaningless.

Premise 2: The question "Is it true that  $X$  is good?" is not meaningless (i.e. it is an open question).

Conclusion:  $X$  is not (analytically equivalent to) good.<sup>4</sup>

Harris calls this argument a "verbal trap"<sup>5</sup> and insists that well-being *is* good *by definition*:

If we define "good" as that which supports well-being, as I will argue we must, the regress initiated by Moore's "open question argument" really does stop. While I agree with Moore that it is reasonable to wonder whether maximizing pleasure in any given instance is "good," it makes no sense at all to ask whether maximizing well-being is "good." It seems clear that what we are really asking when we wonder whether a certain state of pleasure is "good," is whether it is conducive to, or obstructive of, some deeper form of well-being. This question is perfectly coherent; it surely has an answer (whether or not we are in a position to answer it); and yet, it keeps notions of goodness anchored to the experience of sentient beings.<sup>6</sup>

If "it makes no sense at all to ask whether maximizing well-being is 'good'", how do you interact with someone who asks that question? My impression from the book was that if Harris can't convince them that their question is meaningless, he'd have to write them off as just not being worth talking to.

But I don't think questions like *why is well-being good?* are (usually) meaningless or tautological. One thing people often mean is: *can you give me a reason why I should care about well-being?* There

are a variety of conceivable ways to answer this, such as:

1. No, whether you care about well-being or not is an arbitrary choice, or just a fact about your personality determined by your genetics and experiences.
2. Yes, you should want to promote well-being because doing so is actually the best way to accomplish some other goal you already care about.
3. Yes, you should want to promote well-being because according to some convoluted argument vaguely similar to [Kant's or Korsgaard's moral arguments](#), you would be implicating yourself in a contradiction if you did not care about well-being.
4. Yes, you should want to promote well-being, because the desirability of pursuing positive qualia and avoiding negative qualia are [irreducible facts which you can directly perceive](#) when you experience such qualia.

(There may be any number of other answers; sometimes the best answer to a question turns out to be one that we hadn't even realized existed in the space of possible answers until we find ourselves confronted with and persuaded by it.)

Harris's book seems, to me, to be sending mixed messages about whether he accepts something like answer #1 or something like answer #4. On the former interpretation, Harris thinks (almost) everyone just *does in fact* share the same fundamental value (well-being), so we can use that value as common ground for reaching moral decisions, without worrying about how to justify that value. On the latter interpretation, he thinks we can see the correctness of that value by considering extreme cases—almost everyone can see how a life of perpetual suffering and degradation is worse than a life of satisfaction and success<sup>7</sup>.

What frustrates me is that, as discussed above, Harris wants us to define terms in ways that would make the question impossible to even ask. The question is important for a reason he himself articulates well:

Factual beliefs like “water is two parts hydrogen and one part oxygen” and ethical beliefs like “cruelty is wrong” are not expressions of mere preference. To really believe either proposition is also to believe that you have accepted it for legitimate reasons.<sup>8</sup>

If we define “wrong” to mean *reduces well-being*, then the empirical fact that cruelty reduces well-being provides a legitimate reason to believe “cruelty is wrong”. But the obvious next question is whether we have non-arbitrary reasons to avoid reducing well-being in the first place (particularly when it's the well-being of someone else). *This* is the question that makes relativists reluctant to pass judgment, and makes religious people worry about the moral implications of atheism. A satisfying response needs to explain *why* we should care about the well-being of others, not just point out that we naturally *do* care. I think such a response is missing from this book.

Having an account of *why* we should care about well-being might also provide the tools to tackle some of the ambiguities in Harris's theory. For example, he says at one point that we might “not [b]e able to perfectly reconcile the tension between personal and collective well-being”<sup>9</sup>, but I'm not sure his theory provides us any guidance *at all* on that question. How does knowing that *good = promotes well-being* help me decide between pure selfishness, maximizing the total well-being of everyone, maximizing the well-being of the worst-off, maximizing average well-being, etc.? Also, although Harris does discuss the fact the term “well-being” is vague and argues that that's OK, I'm still not clear on how (even in principle, assuming we know all relevant empirical facts) his theory would deal with disputes about the relative importance of different kinds of well-being (e.g. artistic achievement versus physical pleasure).

1. Sam Harris, *The Moral Landscape: How Science Can Determine Human Values*, 1st Free Press hardcover ed (New York: Free Press, 2010), 12.

2. Ibid., 2.
3. Ibid., 62.
4. “Open-Question Argument,” in *Wikipedia*, April 9, 2024, [https://en.wikipedia.org/w/index.php?title=Open-question\\_argument&oldid=1218077466](https://en.wikipedia.org/w/index.php?title=Open-question_argument&oldid=1218077466).
5. Harris, *The Moral Landscape*, 10.
6. Ibid., 12.
7. See *ibid.*, 15.
8. Ibid., 14.
9. Ibid., 188.

*Posted 2024-10-11 by Jacob Williams on [brokensandals.net](https://brokensandals.net). If you have feedback, [email me!](mailto:jacob@brokensandals.net)*