



Book review: Sam Harris' *The Moral Landscape*

Russell Blackford

University of Newcastle, NSW; Editor-in-Chief, *Journal of Evolution and Technology*

russellblackford@bigpond.com

Journal of Evolution and Technology - Vol. 21 Issue 2 – December 2010 - pgs 53-62

The Moral Landscape: How Science Can Determine Human Values. By Sam Harris. Free Press, New York, 2010. 291 pp., \$26.99 (hardcover).

Introduction

In recent years, Sam Harris has become a leading figure in the rational scrutiny of religions and religious cultures, earning himself a place as a prominent “New Atheist,” along with Richard Dawkins, Daniel Dennett, and Christopher Hitchens. To the extent that the New Atheism is a genuine social movement, Harris deserves much of the credit for it. In 2004, he made a dramatic breakthrough when *The End of Faith* was published by W.W. Norton. This was a fiercely anti-religious book, targeted especially at Islam, and emphasizing that religious ideas actually matter because religious adherents are motivated one way or the other to act in accordance with the teachings they accept. The breakthrough was in convincing a major trade publisher to pick up a book like this, and then support it aggressively. Other large publishers followed suit with high-profile critiques of religion by Dawkins and others.

In *The Moral Landscape*, Harris pushes his agenda a step further, examining the nature of morality from a secular viewpoint and offering prescriptions for change. In particular, he contests the moral credentials of religion, argues against popular understandings of free will, and savages moral relativism. He presents an eloquent, passionate, but scholarly defense of his particular take on the phenomenon of morality; he defends moral realism and a consequentialist approach to moral thinking. Harris argues that science can give us the information we need to critique moral systems and develop public policy. If he has his way, much of our moral thinking in the liberal democracies of the West will change quite radically; in particular, we will reject the detached and quietest attitude taken by many Western intellectuals to traditional moral systems. *The Moral Landscape* is an ambitious work that will gladden the hearts, and strengthen the spines, of many secular thinkers.

I enjoyed this book, and I recommend it highly. Though it contains much technical material, from neuroscience as well as philosophy, Harris makes it all accessible. He has an enviable gift for vivid phrasing and clear exposition of difficult concepts, and he undoubtedly has

much to teach us. Almost anyone could benefit from reading *The Moral Landscape*. In that sense, I need go no further. Is this book worth obtaining and reading? Emphatically yes.

Serious reservations about a good book

That said, I have serious reservations. Having now read the book three times, I find that most of the interesting things I could say would be explanations of my concerns and disagreements. Part of the problem, no doubt, is that I would have written a rather different book if I'd tackled the same subject, and of course there is often a temptation for a reviewer to dismiss a book simply for not being what he or she would have written. I'm very conscious of that temptation, and I have no wish to be dismissive, so allow me to emphasize that nothing which follows detracts from *The Moral Landscape*'s obvious strengths or those of its author.

Many passages are very convincing: for example, Harris provides a particularly lucid and penetrating critique of libertarian notions of free will. This alone could stand as an important contribution to public debate, though even here there's room for doubt as to what policy prescriptions should follow. Some of our current policies may be rationalized, by some people, some of the time, on grounds that invoke libertarian free will. Whether they really depend on that idea is another matter. If Harris is right, some policies, such as those relating to criminal justice, may need to be rethought from the ground up; but it remains to be seen how far they'd need to change. In particular, Harris is very quick to dismiss compatibilist accounts of free will and to assume that only libertarian accounts can underpin certain widespread moral intuitions. That remains to be demonstrated.

Be that as it may, Harris sees opponents on two sides. On his right, as it were, are various kinds of moral traditionalists, some of whom support moral ideas that are intellectually untenable, destructive of human happiness, and, in some cases, even cruel. On his left are various kinds of relativists, moral skeptics, error theorists, and non-cognitivists. He seems to think that the latter play into the hands of the former, that their theories suck away our ability to engage in moral critique. That, however, is not necessarily true. It may apply to certain crude kinds of moral relativism, and Harris is impressive in attacking these, but it needn't be so of more sophisticated theories that are taken seriously by professional philosophers. Even if moral error theories, for example, are *disconcerting*, they don't necessarily entail any quietism about tyranny, cruelty, or unjust discrimination.

Unfortunately, Harris sees it as necessary to defend a naïve metaethical position; and, although the defense itself is conducted with considerable sophistication, he does not seem to understand the more sophisticated theories over on his left (or why they are not necessarily in opposition to his main agenda). As a result, I find some of the main lines of argument in *The Moral Landscape* unconvincing, though I accept many of its practical conclusions. In particular, Harris is correct to attack popular, philosophically unsophisticated, forms of moral relativism, and to encourage our hostility to traditional moral systems that cause suffering and harm. Liberal tolerance has its merits, but we'd better make sure we're tolerating the right things, things that are largely harmless.

However, Harris reaches these conclusions only by offering what strikes me as a highly implausible and ultimately unsustainable account of the phenomenon of morality. That account does not seem necessary to reach his practical conclusions, or at least something very like them, but I fear that he'll convince some readers otherwise. We can live with a more sophisticated view of morality than the one Harris offers while getting to a similar place in the end.

The problem lies in his insistence that moral judgments, such as "Lying in circumstances C is morally wrong," are straightforwardly and determinately true or false in the same way as factual statements, such as "My breakfast mug contains coffee," appear to be. We may tend to

think of both kinds of statements in the same way, and it may be unsettling to realize that morality isn't quite like that. If, however, as I'm convinced, it's not, then we'd better try to get understand *how* and *why* it's not, and whether there are any important practical implications. Unfortunately, Harris is impatient with all this, and often resorts to outright scorn in rejecting considerations that don't fit with his position.

The picture according to Harris

Here is how the picture looks if we go along with Harris. Ordinary factual claims are straightforwardly and determinately true or false, as are the theoretical claims made by science. So are moral judgments, and in much the same way. Indeed, moral judgments are simply claims about the well-being of conscious creatures – claims that may often depend on scientific evidence. Of course, Harris acknowledges, we may often face practical difficulties in establishing *whether* a claim about the well-being of conscious creatures is true or false. In principle, however, there will always be an answer. Compare a claim about the number of grains of sand on a particular beach or the number of blades of grass in my neighbor's lawn. In each case, Harris thinks, there is a correct answer, as long as the question itself has sufficient precision. It may be impossibly difficult to ascertain the answer in practice, but we can easily distinguish between answers that are somewhere in the vicinity from those that are not.

Surely there's something about this that sounds attractive. Morality has *something* to do with the well-being of conscious creatures, or so it seems to me. When moral systems lose sight of this, they lose much of their point (don't they?) and are likely to become counterproductive, harsh, or even cruel. Harris does well to point this out and to argue for it at length. It's an important take-home lesson. But as I'll come to, Harris goes much further.

Surely, too, Harris has a point in arguing that science can inform our choices, including those which we label "moral." If our aim is to reduce suffering, for example, science may offer us information about how to do so. As we discover more and more about the world, our developing moral ideas may increasingly be molded by advances in scientific knowledge. Furthermore, Harris is surely correct to deprecate any clear boundary between science and other areas of empirical inquiry, such as the investigative work of historians. He makes the compelling point that rational inquiry into the world around us (and into our own psychological nature) can provide crucial information for practical decision-making. We are still a long way, I suggest, from a situation where we can discard such things as folk understandings of what makes people happy, our own accumulated experiences as individuals, and insights from literature; and we must continue to reflect on all of these things. In principle, however, much useful information can be obtained from more formal kinds of empirical inquiry.

At the same time, however, Harris overreaches when he claims that science can *determine* human values. Indeed, it's not clear how much the book really argues such a thing, despite its provocative subtitle. Harris presupposes that we should be motivated by one very important value, namely the well-being of conscious creatures, but he does not claim that this is a scientific result (or a result from any other field of empirical inquiry). If, however, we combine this fundamental value with knowledge as to how conscious creatures' well-being can actually be aided, we can then decide how to act. We can also criticize existing moral systems, customs, laws, political policies, and so on, if we are informed by scientific knowledge of how they affect the well-being of conscious creatures.

While this is all coherent, Harris is not thereby giving an account of how science can determine our most fundamental values or the totality of our values. If we presuppose the well-being of conscious creatures as a fundamental value, much else may fall into place, but that initial presupposition does not come from science. It is not an empirical finding. Thus,

even if we accept everything else in *The Moral Landscape*, it does not provide an account in which our policies, customs, critiques of policies and customs, and so on, can be determined solely by empirical findings: eventually, empirical investigation runs out, and we must at some point simply presuppose a value at the bottom of the system, a sort of *Grundnorm* that controls everything else.

Harris is highly critical of the claim, associated with Hume, that we cannot derive an “ought” solely from an “is” – without starting with people’s actual values and desires. He is, however, no more successful in deriving “ought” from “is” than anyone else has ever been. The whole intellectual system of *The Moral Landscape* depends on an “ought” being built into its foundations.

“Well-being”

This brings me to an obvious problem in the book, though certainly not the deepest one. For Harris, the key value from which everything else follows is “the well-being of conscious creatures”; however, it’s difficult to know just what is meant by “well-being.” We get the general idea, of course. It is possible to describe situations where somebody, or something, is enjoying well-being – and everyone will agree. It’s possible, too, to describe dramatically different situations where we all recognize suffering and hardship, and we’ll have no difficulty in concluding that the creatures involved are not enjoying well-being at all. Things are going badly for them. Harris offers plausible examples of both classes of situations.

Nonetheless, there are cases where the situation is far less clear. We may find it difficult to judge who enjoys more well-being than whom. Harris is correct to point out that this does not have to be fatal to his approach. After all, many other concepts, such as that of health, are fuzzy around the edges, yet usable in practice. Perhaps “well-being” is like that. We cannot nail the concept down precisely, but we recognize well-being when we see it, and we can promote it without worrying too much about precise definitional niceties. We can also recognize situations where a system of customs or laws is not promoting well-being and may even be harming it.

That’s all fine as far as it goes, and I would have less problem if Harris put it that way consistently. He could insist that the point of moral systems is to protect and promote well-being, while acknowledging that well-being is an inherently fuzzy concept and open to legitimate disagreement at the margins. The concept might then be a place-holder for something else for which it stands as a first approximation: it might be a kind of summation of other things that we value, such as pleasure, satisfaction of preferences, and the possession of various functional capabilities. Unfortunately, that would not assist Harris in insisting that moral questions have determinate, objectively correct answers. There could be situations where the question of which course of action might maximize well-being has no determinate answer, and *not* merely because well-being is difficult to measure in practice but because there is some room for rational disagreement about exactly what it is. If it’s shorthand for the summation of various even deeper values, there could be room for legitimate disagreement on exactly what these are, and certainly on how they are to be weighted. But if that is so, there could end up being legitimate disagreement on what is to be done, with no answer that is objectively binding on all the disagreeing parties.

I suggest that this approach is more plausible than the one taken by Harris himself. Moreover, it need not produce results greatly different from his own in actual practice. Our various conceptions of “well-being” would not all be identical, but they would have considerable similarities, which would allow for much agreement in practical situations. We could reach consensus on many issues, while also reaching a principled understanding of why *total* consensus is not possible.

Harris, however, appears committed to the view that there are determinate and objectively correct answers to all moral questions, even if we cannot discover them in practice. He acknowledges the theoretical possibility that two courses of action, or, say, two different systems of customs and laws could be *equal* in the amount of well-being that they generate. In such cases, the objectively correct and determinate answer to the question of which is morally better would be: "They are equal." However, he is not prepared to accept a situation where two people who have knowledge of all the facts could legitimately disagree on what ought to be done. The closest they could come to that would be agreement that two (or more) courses of action are equally preferable, so either could be pursued with the same moral legitimacy as the other.

According to this picture, well-being is *something that has a metric*. But what is this "something"? While Harris is impatient with what he sees as unimaginative conceptions of well-being, he needs it to be something that is measurable on a scale, so that objective comparisons can be made. When the drift of the argument presses him towards defining well-being, he says that he is not talking about feelings of pleasure; instead, he tends to invoke ideas of deep satisfaction or fulfillment. That seems problematic on its face, because it is far from clear that all conscious creatures are capable of experiencing anything like this. If they are conscious at all, I suppose that they can experience physical pains and pleasures, but how much more than that is experienced by, say, an alligator? We really don't know and should not make assumptions. Presumably the metric for well-being must apply to the most primitive kinds of sensations as well as to psychological satisfactions and feelings of fulfillment.

In the end, I doubt that there really is a metric that we can use to gain fully determinate answers to questions of what will maximize well-being. As I suggested above, however, we don't need this in practice. We can obtain information about such things as physical pleasures and pains, psychological suffering, feelings of fulfillment, someone's objective ability to do things and achieve goals, and so on, without believing that all of these can be reduced to a single metric. The information will still be useful in guiding actions and policies, in criticizing laws and customs, and so on, and it will often guide us to agreement among ourselves. In some cases, most or all of the information will point in the same direction. At a minimum, we should be able to rule out many actions and approaches, and to condemn many existing social arrangements. Some disagreement may remain, but surely we can live with that. After all, we'd need to live with practical disagreements even if well-being *did* have a metric, since even then, we'd often be plagued by the problems of practical measurement.

Are moral judgments objectively binding?

As I've indicated, however, the definition and measurement of well-being is not my deepest objection to the Harris approach. For the sake of argument, and as implausible as it may seem, let's assume that well-being is a single thing that can be specified and quantified. We can assume that, at least in principle, there is always a correct and determinate answer as to what course of action (an individual choice or a political policy, for example, or the continuation or abandonment of a traditional custom) will maximize overall well-being for whatever conscious creatures are affected. Even so, it is not obvious that I should, in any particular set of circumstances, adopt the course of conduct with that effect. Why, for example, should I not prefer my own well-being, or the well-being of the people I love, to overall, or global, well-being? If it comes to that, why should I not prefer some other value altogether, such as the emergence of the *Übermensch*, to the maximization of global well-being?

Thus there might be a determinate, objectively correct answer to what maximizes global well-being, but no such answer to the ancient questions, "How am I to act?" and "How am I to live?" It's *these* questions that really matter, if we're looking for guidance for our actions.

Harris never provides a satisfactory response to this line of thought, and I doubt that one is possible. After all, as he acknowledges, the claim that “We should maximize the global well-being of conscious creatures” is not an empirical finding. So what is it? What in the world makes it true? How does it become binding on me if I don’t accept it?

Sometimes Harris seems to think that the course of conduct which maximizes global well-being is the morally right one because “morally right” just *means* something like “such as to maximize global well-being.” But this won’t do. It’s no use telling somebody (we’ll call her Alice) to act so as to maximize global well-being on the ground that this is the morally right thing to do, while *also* telling her that “morally right” just *means* “such as to maximize global well-being”: the upshot is that Alice is told to act to maximize global well-being because this will maximize global well-being! That’s circular. If she is more committed to a goal such as maximizing her own well-being, or that of her loved ones, than to maximizing global well-being, she is not thereby making a mistake about anything in the world. Nor is she doing anything self-defeating, if she maximizes her own well-being, or that of her loved ones, whenever these conflict with maximizing global well-being.

Nor is she necessarily making any mistake about the world, or doing anything self-defeating, if she chooses to assist the emergence of the *Übermensch*, should that goal clash, in a specific case, with maximizing global well-being. Global well-being is simply not what she most desires or what she prioritizes as her goal. As she might tell us, “End of story!”

Maximizing global well-being, whatever that really amounts to, may have *some* attraction for most or all of us, but someone who cares more about something else, and acts accordingly, does not thereby make a mistake about the world or do anything self-defeating. If she is told that she should adopt course of action X because it will maximize global well-being, she may, quite rationally, point out that she prefers to adopt course of action Y because it will maximize her own well-being and that of her loved ones – and that this is her higher priority. It’s no use telling her that course of action X is the morally right one if it becomes apparent that that just *means* “the course of action that will maximize global well-being.” She can adopt the translation of “morally right” and ask again why she should maximize global being.

Let’s tease this out a bit further. If “morally right” meant something like “what people should actually do, all things considered,” it would, of course, be incoherent for Alice to ask whether she should do the morally right thing. This would translate to the nonsensical, “Why should I do what I should do, all things considered?” However, on the theory we’re considering, “morally right” does *not* mean this. If “morally right” means something like “such as to maximize global well-being” it becomes coherent for Alice to ask why she should do the morally right thing. It translates to the perfectly coherent and understandable question, “Why should I act in such a way as to maximize global well-being?”

At one point, Harris toys with the rather desperate idea that even the word “should,” or the expression “ought to,” can be translated along the lines that “You should do X,” or “You ought to do X” *means* “X will maximize global well-being.” Apart from the inherent implausibility of this for any competent speaker of the English language, it misses the point. Suppose we assigned that meaning to the word “should”: we could then translate the question “Why should I act in such a way as to maximize global well-being?” as the ludicrous, “Does acting in such a way as to maximize global well-being maximize global well-being?” Of course it does, but this gets us nowhere. When Alice proposes to do Y, even though alternative course of action X will maximize global well-being, she is not asking us which of the possible courses of action will, in fact have that effect. She already *knows* this.

Her question is, in effect, “What is that to me?” How, she wants to know, is she rationally required to act against her preference? How, for example, is she acting against a standard that she herself accepts? Or what vital information about the world is she still missing? Or how

will her proposed course of action frustrate her own goals or otherwise be self-defeating? As Richard Joyce argues in detail in his splendid book *The Myth of Morality* (2001), a point can be reached where someone like Alice is not failing to understand anything or acting against her own standards or frustrating her own goals (even if her behavior is socially frowned upon, she may, quite correctly, believe that she can get away with it in this case). Under those circumstances, there is a sense in which the prescription to adopt course of action X is not *rationally binding* on her. She is not failing to *understand* anything. And unless we cheat by invoking a standard of rationality that is already moralized in some way, she is not doing anything *irrational* when she goes ahead with action Y.

If we want to persuade Alice to take action X, we need to appeal to some value (or desire, or hope, or fear, etc. ... but you get the idea) that she actually has. Perhaps we can appeal to her wish for our approval, but that won't work unless she actually cares about whether or not we approve of her. She is not rationally bound to act in the way we wish her to act, which may be the way that maximizes global welfare, unless we can get some kind of grip on her own actual values and desires (etc.).

Harris does not seem to understand this idea. In a long endnote to *The Moral Landscape*, he accuses J.L. Mackie of making an elementary error in developing the idea in his celebrated *Ethics: Inventing Right and Wrong* (1977). Harris says, "Clearly, Mackie has conflated" two notions of objectivity and subjectivity. Leaving aside Harris' habitual over-reliance on the words "clear" and "clearly," often to support assertions that are not clear at all, Mackie makes no such error.

Mackie does not confuse the idea that talk about morality relates to our experience (and is in that sense "subjective") with the idea that it is therefore biased or merely personal ("subjective" in a different sense). The point I've been making, and which Mackie makes throughout *Ethics: Inventing Right and Wrong*, is that claims about how we should act are not rationally binding on us irrespective of such things as our actual values and desires. Putting it another way, they are not *objectively prescriptive*. Alice is never *rationally bound* to act so as to maximize global well-being, if she actually places a higher priority on something else, such as maximizing the well-being of herself or her loved ones. That is what is frequently meant, including by Mackie, when philosophers say that morality is not "objective." As Richard Joyce shows, there is always room for questions such as "What is that to me?"

There can, as Harris states, be *objectively true facts about people's subjective experience*. But Mackie and Joyce are not disputing this. Their point is that there are no judgments about how people like Alice should conduct themselves *that are binding on them as a matter of fact or reason, irrespective of such things as what they actually value, or desire, or care about*.

Of course, Alice might be told to act in such-and-such way – to do X for example – on the basis that it is a requirement of the law, or of the rules of etiquette, or the customs of her society, or perhaps the moralized expectations of people in her sub-culture and social class. There is then a sense in which she is objectively bound to conform to an institutional norm. That is, there will be an objective fact as to whether she has breached the norm or not. But that is not very interesting: even the crudest moral relativists do not deny that people really can breach, or else conform to, institutional norms. But Alice may not be psychologically committed to obeying the institutional norm, and she may be prepared to accept whatever penalty may be imposed for its breach. She may even be in a position to evade it. Perhaps she can be made to conform if she fears the penalty, but even this will require an appeal to her desires: her reasons for conformity will not transcend the social institution and her psychological aversion to punishment.

This is important, because it demonstrates why no one can ever say, coherently, "You *just should* do X." That is not a coherent notion because it is always possible for Alice to ask such

things as what she is still misunderstanding, or what she is doing that is self-defeating, if she goes ahead and instead does Y. If we are going to provide her with reasons to act in a particular way, or to support a particular policy, or condemn a traditional custom – or whatever it might be – sooner or later we will need to appeal to the values, desires, and so on, that she actually has. There are no values that are, mysteriously, objectively binding on us all in the sense I have been discussing. Thus it is futile to argue from a presupposition that we are all rationally bound to act so as to maximize global well-being. It is simply not the case.

Of course, many of us will form an adverse attitude toward someone who does not care at all about the well-being of others. Such a person will probably not be good to have around – indeed, she will be a danger to the rest of us. But she is not necessarily making a mistake about anything. In any event, Alice may not be like that – not on the information provided so far. Perhaps her favored course of action Y will contribute less to global well-being than course of action X, but only by a small amount. If Alice chooses course of action Y, some significant number of people will, let us say, end up less happy than might have been the case if she'd chosen action X. But perhaps action X, while providing a slightly greater amount of (net) global well-being than course of action Y, will produce considerable pain and suffering for one of Alice's children (and thus for Alice herself). In those circumstances, Alice's course of action may be quite *rational*, given her goals, but, more than that, we are also likely to assess it as *reasonable* and acceptable to us, even if we would still prefer that Alice adopt course of action X.

In a circumstance such as this, the well-being of Alice and her child is, as it were, in competition with that of certain others. However, we usually accept that people act in competition with each other, each seeking the outcome that most benefits them and their loved ones. We don't demand that everyone agree to accept whatever course will maximize the well-being of conscious creatures overall. Nothing like that is part of our ordinary idea of what it is to behave morally. By everyday standards, Alice is not beyond the pale; she is not someone not worth talking to and socializing with. By everyday standards, she is not an immoral person, and course of action Y is not morally forbidden.

Moral systems and their demands

This brings me to the important point that actual systems of moral norms do not make superlative demands such as the demand that we act to maximize the global well-being (whatever that is, exactly) of conscious creatures. Rather, they permit us to act as we wish within certain boundaries that may differ considerably, from system to system, in how restrictive they are. Some actions will be compulsory, others will be forbidden, but the remainder are left to the discretion of the individual concerned. There may also be actions that are praised or frowned upon, but not required or forbidden. And there may be more complex ideas of good character, relating to the dispositions of individuals to act "well" or "badly" in the various circumstances that arise. Some moral systems are very prescriptive indeed, but they all leave at least some room for competition, conflict, and individual discretion.

Though it is difficult to prove to a high degree of certainty one way or the other, I suggest that these systems do not, in fact function (indirectly and ineffectually) to maximize global well-being. They are likely to have less glorious, more down-to-earth social functions, and they permit biases towards ourselves and loved ones. Even when moral norms and systems seem aimed at rather broad goals, such as ameliorating suffering, these goals are likely to be limited. It is one thing to follow a norm of avoiding unnecessary suffering, or taking steps to reduce it; it's quite another to set aside one's personal priorities and biases, and to aim at some kind of global maximization of fulfillment or satisfaction.

In the upshot, we can ask what is the *point* of morality and not expect a single, fully determinate answer. We can say with confidence that moral systems do not have so ambitious

a point as the maximization of global well-being, but it is difficult to be specific as to what they do or ought to aim at, short of this. They seem to relate to a number of things that tend to matter to human beings, given our nature as social animals with an evolved psychology (and with certain abilities and limitations). Moral systems, and associated bodies of custom and law, provide us a degree of personal security, reduce suffering, enhance the prospect of (at least intra-societal) peace, and assist social cooperation. They may go further and aim more broadly at amelioration of suffering. Some of the ways all this is done can be seen across many societies and cultures: methods of allocating property, restrictions on the use of violence and fraud, and so on.

That is not to say that *just anything* can be the point of moral system, or that it's all just arbitrary. Given our nature as social animals with certain abilities, limits, propensities, and physical needs, it is unsurprising if the contents of actual moral systems have commonalities. Similar sorts of restrictions on human conduct (and similar requirements or obligations) are likely to be needed in many sorts of physical and cultural environments, and in many economic and technological circumstances. However, moral systems, and the values by which we assess them, can be *non-arbitrary* without those values being fully determined by an objective reality, independent of people's actual values and desires.

To make this point is not to deny that there is, in principle, an objectively true description of the nature of morality as a familiar social phenomenon. On the contrary, I have just given what I believe to be the beginnings of such a description. There are truths about the phenomenon of morality, just as there are truths about other things, but one of the truths about morality is that there's a sense in which we are not objectively bound by moral norms. I expect that most educated people have a suspicion that something about morality is "subjective" or "relative" or "not objective," even if they can't quite put their fingers on it. And yes, there is something to the suspicion, but when we nail down what that actually is we need not be alarmed. It does not make morality *just arbitrary* or *capable of taking any form*. It doesn't prevent us developing coherent, rational critiques of various systems of laws or customs or moral rules, or persuading others to adopt our critiques.

All this may be disconcerting, since it overthrows naïve ideas of morality. It may even mean that ordinary first-order moral claims are literally false, to the extent that they assert (or are heard to assert) their own objective prescriptivity or bindingness. My analysis will meet much resistance, emotional as much as intellectual, but why, on deeper reflection, is it terribly surprising? Many of our value judgments are non-arbitrary without involving fully determinate, objectively correct answers. No one expects *that* kind of answer to a bald question such as, "What is the best refrigerator?" or "What is the best motor-car?" Our criteria for judging the merits of refrigerators and motor-cars are not arbitrary, but nor are they reducible to a metric that we're all objectively bound to employ. Judgments about refrigerators and motor-cars can be rational and defensible, but there is also room for legitimate disagreement. As manufacturers know, we all want very similar things from motor-cars (or refrigerators) but we needn't all want *exactly the same* things.

Something like this applies to judgments of the merits of Victorian novels, 1950s science fiction movies, stage actors, paintings and sculptures, houses and gardens, friends and lovers, computer firms, mathematics teachers, philosophers, pet cats and dogs, sunsets, public speakers, winter holidays, rose varieties, pop songs, academic journals, cake recipes, photographs of Saturn's rings, and most other things whose merits concern us at all.

Conclusion

Despite my criticisms, Harris is correct on the most important point. We can criticize other cultures as well as our own. Popular moral relativism notwithstanding, we need not adopt a quietism about moral traditions that cause hardship and suffering. Nor need we passively

accept the moral norms of our own respective societies, to the extent that they are ineffective or counterproductive or simply unnecessary.

In particular, it is quite open to us to condemn traditional systems of morality to the extent that they are harsh or cruel, rather than providing what most of us (quite rationally) want from a moral tradition: for example that it ameliorate suffering, regulate conflict, and provide personal security and social cooperation, yet allow individuals a substantial degree of discretion to live their lives as they wish. We don't all have to agree on exactly what weight to give to these, or whatever other things of a similar kind we might want from a moral tradition. Even without total agreement on the *point* of morality and on *exactly* what we want from a moral tradition, we are quite capable of converging on similar judgments in many real-world cases. We can, for example, agree to repudiate, and speak out against, any system of norms that suppresses female sexuality – and women's freedom more generally – through genital mutilation, veiling, and social isolation.

More generally, we can step back and ask whether a society's rules, customs, and laws are actually contributing to the sorts of outcomes that we can (non-arbitrarily) want from them. At the other extreme, we can ask whether they are producing frustration, hardship, and suffering.

Unfortunately, many societies rationalize their moral traditions on a false basis, such as portraying their moral rules as the commands of a deity. Many people may come to believe that obedience to the supposed deity is the real point of the rules that they follow and seek to impose on others. Convincing them otherwise may be impossible without changing their comprehensive understandings of the world, which may, in turn, not be practicable if they have been thoroughly socialized into a particular viewpoint from childhood, or if they have a deep emotional investment in it. But why should *that* be surprising? Even if Harris were correct in his strongly realist account of the phenomenon of morality, he would face a pressing difficulty: that of getting people with deeply entrenched worldviews that are radically different from his to adopt his broadly utilitarian account of normative ethics ... without first abandoning their religious or metaphysical views.

In the end, Harris provides a compelling argument for selective intolerance toward harsh moral traditions. He argues via a kind of moral realism, linked to a form of utilitarian ethic, but I submit that these are not doing the real work. To reach a similar conclusion, we can rely on much weaker premises. It's enough that we have a non-arbitrary conception of what morality is for, and what sorts of things we can rationally and realistically want moral traditions to do. Where they divert from that conception, moral traditions merit our critique and opposition. These should be every bit as severe, absolutely as passionate, as Harris evidently wants, but that does not commit us to his total picture of morality's landscape.

Like it or not, morality is a much trickier phenomenon.

References

Harris, S. 2010. *The moral landscape: How science can determine human values*. New York: Free Press.

Joyce, R. 2001. *The myth of morality*. Cambridge: Cambridge University Press.

Mackie, J.L. 1977. *Ethics: Inventing right and wrong*. London: Penguin.