

UC Berkeley

ISUS-X, Tenth Conference of the International Society for Utilitarian Studies

Title

"Reconciling Consequentialism with Ordinary Moral Knowledge"

Permalink

<https://escholarship.org/uc/item/58r6d1bs>

Author

Lukas, Mark

Publication Date

2008-09-11

RECONCILING CONSEQUENTIALISM WITH ORDINARY MORAL KNOWLEDGE

Mark Lukas
Longwood University

Abstract

According to objective consequentialism (OC), an act is morally right iff no alternative has better consequences. Some object to OC on epistemic grounds: because we cannot foresee the long-term consequences of our alternatives, let alone accurately predict their relative values, OC is useless as a guide to morally right action. A common response to such objections involves construing OC as a criterion of moral rightness rather than a practical action-guiding principle. The idea is to concede that OC is largely useless as a practical principle but maintain that it is nevertheless true. A stronger but largely neglected epistemic objection runs as follows: if OC is true, then because we can never know which of our myriad alternatives would have the best consequences, we never know that any particular act is morally right; but we frequently do know that particular acts are right, so OC is false. In this paper I develop and respond to the stronger objection. One key response plays on the idea that our moral beliefs about particular acts are often (defeasibly) justified by moral beliefs about relevant act types, beliefs that are in turn often justified by rational moral intuition.

INTRODUCTION

According to objective consequentialism (OC), an act is morally right iff no alternative has better consequences. One familiar brand of OC is Act Utilitarianism (AU), the view that an act is morally right iff it maximizes utility. Lamentations about the impracticality of both views, as well as related utilitarian doctrines, flood the literature. Such complaints generally play on our epistemic limitations regarding our alternatives and their consequences. But these complaints are nothing new; utilitarians have been aware of them for at least a century and a half now. Mill, Sidgwick, and Moore all acknowledged the obvious fact that it is often very difficult if not impossible to ascertain

whether a given act satisfies the utilitarian standard for rightness, and so it may very well be extremely difficult to "apply" the standard as a practical guide to action.¹

About forty years ago, in an especially enjoyable paper, Eugene Bales pointed out that even if we concede the point that act utilitarianism is useless, or nearly useless, as a practical guide to action, utilitarians need not give up on their theory, for it might nevertheless be true.² It might be true that an act is morally right iff it maximizes utility and yet be impossible in practice to determine which acts actually maximize utility.

Notwithstanding Bales' paper, epistemic objections persist. Shelly Kagan rehashes the old chestnut as follows: "it is impossible to know the future....This means that you will never be absolutely certain as to what all the consequences of your act will be ... there may be long term bad effects from your act, side effects that were unforeseen and indeed unforeseeable...Yet if it is impossible to tell whether any act is morally right or wrong, how can consequentialism possibly be a correct moral theory."³ Robert Frazier argues that if AU is true, then "we cannot be justified in believing of any act that performing it is morally permissible."⁴ The obvious implication here is we can be thus justified, at least sometimes, and so AU is false. More recently, James Lenman has argued that since we are all essentially clueless about the consequences of our actions consequentialism, though perhaps true as Bales suggests, is not very interesting since it has "nothing whatsoever to do with our practical thought in its application to the actual world."⁵

In what follows I defend OC by defending its most prominent manifestation: AU. I defend AU from an argument that is akin to the epistemic objections I've been alluding

¹ See, for example, John Stuart Mill's *Utilitarianism*, ed. John M. Robson, Toronto, *Collected Works of John Stuart Mill*, (1965), x, 224. Also see, for example, Henry Sidgwick's *The Methods of Ethics*, The University of Chicago Press, (1962) Chapter IV. See G.E. Moore's *Principia Ethica*, Cambridge, (1903), Chap. V, Sections 90 & 91

² Eugene Bales, "Act-Utilitarianism; Account of Right-Making Characteristics or Decision-Making Procedure?", *American Philosophical Quarterly*, Vol. 8, Issue 3, (1971), 257-65.

³ Shelly Kagan, *Normative Ethics*, Boulder, (1998), 64.

⁴ Robert Frazier, "Act Utilitarianism and Decision Procedures", *Utilitas*, Vol., 6, Issue 1, (1994), p.44

⁵ James Lenman. "Consequentialism and Cluelessness", *Philosophy and Public Affairs*, Vol. 29, No. 4, (Autumn, 2000), 342-70

to. The argument is rarely discussed but it lurks in the shadows. The argument most resembles the main argument offered by Robert Frazier in his paper, and it is a direct challenge to Eugene Bales suggestion that the truth of utilitarianism is consistent with our near total ignorance regarding our alternatives and their consequences. After some preliminary stage setting, I present and explain the argument. I then discuss several responses to the argument, one flawed response and several more promising responses. My goal is to show that our epistemic situation regarding what is right and what is wrong is not nearly as bleak as objectors to OC and AU would have us believe. For I believe both that OC and AU are true and that when it comes to a great many quotidian acts, the acts that each of us perform in the thousands, or perhaps millions, everyday we often do know which are right and which are wrong.

STAGE SETTING

Prior to our engaging in any rigorous philosophizing about morality, most of us take it for granted that we know quite a bit about which acts are morally right and which are wrong. Most of us feel quit certain that Hitler performed some morally wrong acts and that Mother Theresa performed a great many acts that were morally permissible, if not supererogatory. But even with respect to many more mundane acts performed by ordinary people, most of us are usually quite confident in our beliefs about which are right and which are wrong. I am now scratching my ear, and I'm pretty sure that in doing so I'm not doing anything morally wrong. Indeed, I'm willing to assert that I know that my ear-scratching is morally permissible. I'm also sure that if I were to blow up a school bus full of little children I would thereby perform an act that is morally wrong. But if utilitarianism is true, how can I be so sure of these things? Although is seems innocent enough, perhaps scratching my ear will in fact trigger some chain of events that results in mind-boggling amounts of suffering. This is a possibility. Maybe my scratching will outrage one of you; it will drive you insane, and you'll go on a killing rampage. Could happen. And given that I can't be sure this won't happen, how can I be so sure that my act is morally right? Maybe blowing up the bus will prevent the birth of some future Hitler-like character and thereby prevent a cataclysmic world war. This too is a possibility.

Indeed, maybe it is actually my moral obligation to blow up the bus since doing so would be the only way to maximize utility.

Consider a small boy, Harry. Harry is ten years old and he is playing in his own backyard. As I observe Harry from a distance, I notice that he's pretending to play soccer. He is alone, but he has imagined himself to be surrounded by other players and adoring fans. He has the ball in the final seconds of the World-Cup Final. If he scores, his team wins. Otherwise they lose. In reality, Harry doesn't even have a ball. He has built a small mound of dried leaves and he is pretending that it is a ball. He takes his time eyeing up an imaginary goalie; then he runs forward and kicks the pile of leaves, sending an imaginary ball into an imaginary goal, leading his team to victory. Call Harry's act of kicking the leaves 'A1'. I'm pretty sure that A1 is not morally wrong. In fact, I will say with confidence that I know A1 is morally permissible.

THE ARGUMENT

The epistemic argument against utilitarianism that I want to discuss goes like this:

1. If AU is true, then I do not know that A1 is morally permissible.
2. I do know that A1 is morally permissible.
3. Therefore, AU is false.

Let me explain the rationales behind the premises. First premise one. According to AU an act is morally permissible iff it maximizes utility. An act maximizes utility provided that no alternative has a greater utility. So if AU is true and A1 is morally permissible, Harry has no alternative to A1 with greater utility. That is, there is absolutely nothing he could do that would have better consequences than A1.. But surely, at the time he kicks the leaves Harry has myriad alternatives. Rather than kick the leaves with his right foot, as he does, he could kick them with his left foot. Or he could pause a bit longer and then kick the leaves. Or he could refrain from kicking the leaves altogether and go play with legos instead. Or he could run inside and study German. Maybe he could haul out his microscope and study the leaves up close. Etc..etc... Harry has a great many alternatives.

Surely too many to even grasp with the mind, let alone list here. Moreover, each of Harry's many alternatives has consequences. Some short-term; some long-term. Many of these consequences are unforeseeable. But each alternative has a precise utility determined by its consequences. Given that I can't even grasp, let alone enumerate, all of Harry's alternatives, and given that the precise utilities of these alternatives are a complete mystery to me, how on earth can I be justified in believing that A1 maximizes utility? I can't be, says the objector. And since knowledge requires, at a bare minimum, *justified* true belief, I do not know that A1 is morally permissible.

Now the rationale for premise two. Knowing that A1 is morally permissible is like knowing that I have hands. It is not the kind of thing that any sane person should doubt. It is simply obvious that Harry isn't doing anything wrong when he kicks his imaginary soccer ball. Many mundane acts, like A1, serve as paradigm examples of morally permissible behavior. It is obviously morally ok for me to scratch my ear. And it is obviously morally permissible for me to wave my hand, as I am now doing. Acts such as these provide the raw material for moral theorizing. We notice that a great many acts are obviously morally permissible and that others are obviously morally wrong. Then, if we are philosophically inclined, we might begin to wonder what the right acts have in common and what the wrong acts have in common. We might then cobble together generalizations designed to fit with our common-sense beliefs about particular cases. If we are lucky, we stumble upon the truth. Doubting my knowledge that A1 is morally right smacks of the kind of radical skepticism that has led to some of the most absurd conclusions in the history of philosophy. Clearly I know that A1 is morally permissible. This concludes my presentation and explanation of the argument.

RESPONDING TO THE ARGUMENT

One response to the argument is to bite the bullet and deny premise two. Fred Feldman, a utilitarian, suggests he would make this move. He writes: "I have to acknowledge that if my view is correct, then most of us, most of the time, really don't know what we ought to do...This is indeed unfortunate for us... [But] if the theory represented the epistemology of ethics as being any easier, that would show the theory to

be wrong."⁶ G.E. More says some things to suggest that he too would be sympathetic to this kind of response.⁷ But I find this line of response to be an embarrassment to utilitarians. If a philosophical theory about moral permissibility implies that we can never know of any particular act that it is morally right or wrong, that is reason enough to reject the theory. It's as if I were to present you with a theory purporting to give necessary and sufficient conditions for an animal's being a zebra; but I then tell you that if my theory is correct, we can never actually tell of any given animal whether or not it is a zebra. Such a theory would be ridiculous. I know a zebra when I see one, most of the time. Likewise, I know of a great many particular acts that they are morally permissible; and of others that they aren't. And if AU is incompatible with this knowledge, then as far as I'm concerned AU will have to go.

I think more promising lines of response to the argument involve the denial of premise one. Such responses would allow, us to affirm premise two, as I think we should, for I believe that AU is true and that I know A1 is morally permissible, despite my epistemic limitations regarding Harry's alternatives and the values of their consequences. But how is such knowledge possible? More specifically, what justifies my belief that Harry's act is morally permissible? Let me now sketch two answers to this question.

One way to deny premise one involves adopting a kind of foundationalist epistemology with respect to certain moral beliefs. On this kind of view, some moral beliefs are so obviously true that they are in effect self-justifying. The belief that A1 is morally permissible would be one example. That it is morally ok for me to scratch my nose and wave my hand would be two others. Taking this kind of position relieves the utilitarian of the burden of explaining how we can know the moral status of an act without knowing much about its consequences or the consequences of its alternatives. But is such a foundationalism compatible with utilitarianism? It might be argued that such a view would turn us in to soothsayers of a sort, at least with respect to some acts. If AU is true and, and my belief that A1 is morally right is self-justified, then if A1 is in fact morally right, do I not thereby acquire justification to believe that A1 does maximize utility? It would seem that I have somehow magically acquired justification to believe

⁶ Fred Feldman, *Doing the Best we Can*, D. Reidel Publishing Company, (1986), 48.

⁷ G.E. Moore. *Principia Ethica*, Cambridge, (1903), Chap. V, Sections 90 & 91.

that A1 maximizes utility, and I have acquired this without knowing, or indeed perhaps even believing, anything about its actual utility and the utilities of its alternatives. This all seems a bit mysterious.

But perhaps the mystery is just apparent. I believe that I have hands, and this belief seems self-justified. And it is true that I have hands. Now if some statement to the effect that *I have hands iff x, y, and z* is true, then it would seem that I am justified in believing x, y, and z, though perhaps I don't believe it because I am unaware that these are the necessary and sufficient conditions for my having hands. Here the foundational belief that I have hands, if true, would serve as a kind of indicator that x, y, and z obtain, whether I realize it or not. Likewise, perhaps foundational beliefs about the morality of certain acts serve as indicators that these acts maximize utility, even though we may be unaware of this fact. I confess that I am unsure of the plausibility of this kind of objection to premise one, though I think there might be something to it.

Let me suggest a second way a utilitarian might explain what justifies my belief that A1 is morally right, even though I am almost completely ignorant of A1's utility and the utilities of its alternatives. This response also allows us to hold on to the claim that I know A1 to be right. The explanation involves adopting a kind of hybrid coherentist/externalist epistemology about the justification of beliefs about certain acts. So here's the explanation. When I form my belief that A1 is morally permissible, maybe something like the following is going on. Perhaps very quickly and perhaps only at a subconscious or semi-subconscious level, I notice various features of Harry's act. Among other things, I notice that it is solitary -- although I am observing, I am at a great distance; Harry is alone. And so I perhaps infer that it is very unlikely that his act will have any appreciable effects whatsoever on anyone, including Harry. It seems that his act won't harm anyone. I notice that Harry takes pleasure in performing the act. I notice that the act is an instance of creative, imaginative play. I notice other obvious features of the act and perhaps infer from the fact that he is a child and engaged in typical child-like play that the act is what we might call 'innocent'; i.e., it is not performed with wicked intent. Indeed the act seems not to exemplify any of the vices I'm aware of; e.g., it is not an instance of cowardly, cruel, dishonest, treacherous, malicious, greedy, or gluttonous behavior. And drawing on a lifetime of experience, I realize that acts like Harry's rarely,

if ever, warrant punishment for their agents. Indeed I realize that acts such as his are often encouraged. Goofing around in the backyard is something kids are supposed to do. I realize that I am not aware of a single case where an act such as Harry's has been condemned by anyone as morally wrong. I notice other obvious features of Harry's act as well.

Then my moral intuitions kick in. These intuitions are simply initial intellectual appearances. They are the way certain things initially seem to me, and they are subject to revision. But absent reason to revise or reject them, they provide me with *prima facie* reason to think that certain moral claims are true and others are false. In the present case, here are some relevant moral intuitions: in itself, pleasurable activity is better than painful activity. Non-vicious acts are better than vicious acts. Innocent childhood play is good. Other things being equal, acts that cause harm are worse than acts that don't.

When these, and perhaps other, intuitions get together with the other beliefs about A1 I've mentioned, the gears in my head churn a little and I reach the conclusion that Harry's act is morally permissible. Of course this conclusion is based on limited information. It is based just on the features of Harry's act that I have directly observed or that I have inferred to exist. And it is based on beliefs and moral intuitions all of which are subject to revision should more information arise. Nevertheless, it is in virtue of such a process, I think, that I form my belief and through which my belief gets its justification.

Notice that in telling my story, I didn't run through a list of Harry's alternatives; I didn't mention their relative utilities, and I said nothing to address the question as to whether A1 maximizes utility. I neglected to mention these things because believing that A1 is morally right is not the same as believing that it maximizes utility, even if AU is true. We need to keep in mind that the following two propositions are distinct:

P1. A1 is morally right

P2. A1 maximizes utility

I claim to know P1. I do not claim to know P2. But to know P1, I must be justified in believing it. And if AU is true, does it not follow that to be justified in believing P1 I

must justifiably believe P2?⁸ It does not follow. A great many people have fully justified beliefs about water without any justified beliefs about the relationship between the hydrogen atoms and the oxygen atoms that compose water?" The mere fact that the truth of one proposition is a necessary and sufficient condition for the truth of another does not entail that being justified in believing one of the propositions requires having justified belief in the other. But perhaps what justifies belief in one of the propositions must at least serve as an indicator that the other is, or is likely to be, true, whether the believer in question realizes it or not. In the present case, perhaps what justifies my belief that P1 is true must serve as an indicator of some kind that P2 is true or is likely to be true.

I think that what justifies my belief in P1 does in fact serve as such an indicator. Given my observations of Harry, the beliefs I've formed about the act-types that A1 instantiates, and my moral intuitions, it is reasonable to conclude, I think, that in performing A1 Harry really is doing the best he can, or that it is likely he is. Similar conclusions are warranted regarding my ear-scratching and hand waving.

Of course many philosophers will disagree. They will say that surly Harry and I have better alternatives available. In his paper "Monkeys, Typewriters, and Objective Consequentialism," Eric Willand, says things to suggest I, and perhaps even Harry, both have the physical and intellectual ability to write the mix of marks that, taken together, would compose the Great American Novel, or the journal article that explains how to cure AIDS⁹. Of course we might not realize which particular concatenation of letters, punctuation marks, and spaces would add up to such documents. But there is no doubt in Willand's mind that anyone who knows how to write English can write the relevant mix of marks. And surely, if instead of kicking his leaves Harry were to run in and write an article describing how to cure AIDS, that would have much higher utility than kicking his silly leaves. And if an open alternative to scratching my ear really is to jot down the Great American Novel, then surely I ought not scratch my ear; I should get down to jotting immediately.

⁸ Robert Frazier seems to think so.

⁹ Eric Willand, "Monkeys, Typewriter, and Objective Consequentialism", *Ratio* XVIII (2005), 353-60.

I think Willard and philosophers like him wildly overestimate the alternatives that are open to most ordinary people most of the time. While it is perhaps logically possible and even physically possible for Harry to write down the cure for AIDS rather than kick his leaves, and so in some sense of 'can' he can do it. In a more obviously relevant sense of 'can', he can't do it. Given his current circumstances, which include his current beliefs, desires and other psychological states, writing down the cure for AIDS is simply not one of his open alternatives. Given who he is, where he is, what he believes and desires, it seems to me that Harry's actual alternatives, although perhaps numerous, include almost exclusively mundane acts of the sort one might expect a ten year old boy to perform. These are the acts that are genuine alternatives for him. And I don't think there is any compelling reason to think that among these alternatives is one with better consequences than A1. But of course there might be. I can't discount that possibility completely. But in that case, although my belief in P1 would be fully justified in the sense required for knowledge, it would simply be false. And so I wouldn't know that A1 is morally right. But fallibilism is a fact of life. I don't claim to have *infallible* knowledge of P1. I simply claim to have ordinary, run-of-the-mill knowledge that it's true. And such knowledge is, of course, supported by defeasible justification. And it is just this sort of justification that I claim to have for my beliefs that A1 is morally right and that it is morally ok for me to scratch my ear and to wave my hand.

CONCLUSION

My goal has been to briefly sketch a few ways that AU and other forms of OC might be reconciled with the kind of ordinary moral knowledge that most of us take for granted. I have not tried to show that such views provide us with a reliable practical guide to morally right action. Nor have I tried to show that we can always tell which particular acts are right and which acts are wrong. Moral dilemmas remain. There are obviously a great number of situations in which we simply don't know which is the right thing to do. But when it comes to the overwhelming majority of commonplace acts that each of us perform everyday -- such as ear-scratching, tooth-brushing, limb movement, chewing of food -- we often do know which are right and which are wrong. It simply does not follow

from OC and our (contingent) epistemic limitations regarding our alternatives and their consequences that we only rarely if ever know what is morally right or wrong.