

WORLD UTILITARIANISM AND ACTUALISM VS. POSSIBILISM

I. Hedonistic Act Utilitarianism: Some Deontic Puzzles

Hedonistic Act Utilitarianism (HAU): S's performing x at t_1 is morally permissible if and only if S's doing so maximizes hedonic utility, and if S's performing x at t_1 is morally permissible, then what makes it so is ultimately the fact that S's doing so maximizes hedonic utility.

Hedonic Utility: The hedonic utility of an act equals the total amount of hedons it produces minus the total amount of dolors it produces, where a *hedon* is the standard unit of measurement of pleasure and a *dolor* is the standard unit of measurement of pain.

An act *maximizes hedonic utility* if and only if there is no alternative act available to the agent that would produce more hedonic utility than it would. Note, then, that more than one act can maximize hedonic utility.

A. A Puzzle about Prerequisites

Here's a plausible deontic principle:

The Ought-Extends-to-Prerequisite-Acts Principle: If S ought to do x and cannot do x without doing y , then S ought to do y . More formally:

$$[O(x) \ \& \ \sim C(x \ \& \ \sim y)] \rightarrow O(y).$$

Here, ' $O(x)$ ' stands for 'S ought to do x ' and ' $C(x)$ ' stands for 'S can do x '.

Curiously, hedonistic act utilitarianism violates this deontic principle. To see this, consider the following scenario, which I borrow from Feldman (1986, 5-7):

Suppose that S is, at t_1 , considering what she should do this morning. S could, at t_7 , perform any one of the following mutually exclusive options (and assume that all the other options have much lower hedonic utilities):

a ₁	work in the garden	U(a ₁) = +12
a ₂	go to the dump	U(a ₂) = +8
a ₃	paint the house	U(a ₃) = -15

According to HAU, S should perform a₁ at t_7 , since that is what would produce more utility than any other available alternative.

Now S also has to decide what to do at t_3 . Let's suppose that S could, at t_3 , perform any of the following mutually exclusive options (and, again, assume that all the other options have much lower hedonic utilities):

a ₄	gather garden tools	U(a ₄) = -1
a ₅	load and start truck	U(a ₅) = +1
a ₆	mix paint	U(a ₆) = -2

According to HAU, S should perform a₅ at t_3 . But S can't perform a₁ at t_7 unless she instead performs a₄ at t_3 . So HAU violates the ought-extends-to-prerequisite-acts principle.

As Feldman (1986, 7) suggests, someone might think that amalgamation offers a possible solution to this problem, for consider the following alternatives:

a ₇	gather garden tools and then work in garden	U(a ₇) = +11
a ₈	load and start truck and then go to dump	U(a ₈) = +9
a ₉	mix paint and then paint house	U(a ₉) = -17

HAU implies that S ought to perform a₇. But, far from offering a solution, this only demonstrates that things are even worse than we initially thought, for this example makes clear that HAU not only violates the ought-extends-to-prerequisite-acts principle, but also:

The Ought-Distributes-through-Conjunction Principle: If S ought to do both x and y , then S both ought to do x and ought to do y . More formally:

$$O(x \ \& \ y) \rightarrow O(x) \ \& \ O(y)$$

HAU violates this principle because performing a₇ is just amounts to performing both a₁ and a₄.

B. A Puzzle about Time

Consider the following puzzle, as presented by Feldman (1986, 11). "It involves a doctor and a patient. The patient has some terrible disease. There are two ways in which the doctor can cure the disease. He can give two successive doses of medicine A, or he can give two successive doses of medicine B. Two doses of A would produce a speedy, safe cure. Two doses of B would produce a somewhat less satisfactory cure. And mixture of medicines, or failure to administer medicine, will be fatal. Let's agree that the utilities of these various courses of treatment are as catalogued" below:

a ₁	A on Monday, A on Tuesday	U(a ₁) = +100
a ₂	B on Monday, B on Tuesday	U(a ₂) = +75
a ₃	A on Monday, B on Tuesday	U(a ₃) = -100
a ₄	B on Monday, A on Tuesday	U(a ₄) = -100
a ₅	anything else	U(a ₅) = -100

Now assume that the doctor is going to administer B on Monday even though he could quite easily administer A on Monday instead. Also, assume that it takes all day to administer these medicines and thus that it is impossible to administer both A and B on any given day. Given these assumptions, what should we say that the doctor ought to administer on Tuesday? On the one hand, it seems clear that the doctor ought to administer A on Tuesday, for it is clear that he ought to administer A on Monday and again on Tuesday. This is, after all, the best he can do. On the other hand, if he's going to administer B on Monday, then it seems clear that he ought to administer B again on Tuesday, for doing anything else would kill the patient. So it seems that a good case can be made for each of the following two conclusions: (1) the doctor ought to administer A on Tuesday and (2) the doctor ought to administer B on Tuesday. But the doctor can't do both. Thus, we seem to have here a violation of another plausible deontic principle:

A Corollary of the Ought-Implies-Can Principle: If S both ought to do x and S ought to do y , then S can do both x and y . More formally:

$$[O(x) \ \& \ O(y)] \rightarrow C(x \ \& \ y)$$

Clearly, then, we need to find some solution to these puzzles. That is, we need to find some way of formulating utilitarianism so that it doesn't violate any of these plausible deontic principles.

II. A Solution to These Deontic Puzzles: Future-Course-of-Action Utilitarianism

A. The Theory: Future-Course-of-Action Utilitarianism

FCAU It is, as of t_n , morally permissible for S to perform x at t_m ($m > n$) if and only if there is a utility-maximizing FCA available to S at t_n in which S performs x at t_m .

On future-course-of-action utilitarianism, an act is morally permissible if and only if it is part of some future course of action that would maximize utility.

B. Some Clarifications

1. A Future Course of Action

At any given moment, t_n , prior to death, there are various ways a subject, S, might live out the rest of her life. For each possible way of completing S's life, there is a whole series of actions that S will perform if and only if S completes her life in exactly this way. Each of these possible ways of acting over the remainder of S's life is what I call a future course of action—an "FCA" for short—available to S at t_n .

2. A Utility-Maximizing FCA

An FCA is utility-maximizing if and only if there is no available alternative FCA that would produce more utility than it would.

3. Availability

To say that some FCA is available to S at t_n is to say that, as of t_n , S has both the ability and opportunity to perform the sequence of acts of which that FCA is composed. Note that, if compatibilism is true, then S can, as of t_n , have both the ability and opportunity to perform x at t_m even though S is causally determined to perform y at t_m . Note also that the number of FCAs available to S diminishes with each passing moment.

C. How FCA Utilitarianism Solves Certain Deontic Puzzles

A. The Puzzle about Prerequisites

Assume that the only utility-maximizing FCA available to S at t_1 is one in which S performs a_4 at t_3 and performs a_1 at t_7 . Unlike HAU, then, FCAU implies that it would be wrong for S to perform a_5 at t_3 despite the fact that a_5 would produce more utility than any other available act-alternative. Thus FCAU does not violate either the ought-extends-to-prerequisite-acts principle or the ought-distributes-through-conjunction principle.

B. The Puzzle about Time

Which medicine should the doctor administer on Tuesday? According to FCAU, the answer depends on what day it is, for, on FCAU, obligations are time-relativized. For instance, FCAU implies that, as of Sunday, the doctor ought to administer A on Monday and then again on Tuesday, for the best future course of action available to the doctor as of Sunday is the one where the doctor administers A on both Monday and Tuesday. Thus, as of Sunday, the doctor ought to administer A on Tuesday. But, as of Monday

(after already having administered B on that day), the best future course of action available to the doctor is the one where the doctor administers B on Tuesday. As of Monday, then, the doctor ought to administer B on Tuesday. So FCAU does not imply the following two incompatible conclusions: (1) the doctor ought to administer A on Tuesday and (2) the doctor ought to administer B on Tuesday. Instead, FCAU implies the following two compatible conclusions: (1) the doctor ought, as of Sunday, to administer A on Tuesday and (2) the doctor ought, as of Monday, to administer B on Tuesday. And these verdicts do not conflict with the following ought-implies-can principle, suitably revised in light of the fact that obligations are time-relativized:

A Corollary of the Ought-Implies-Can Principle: If S both ought, as of t , to do x at t' and S ought, as of t , to do y at t'' , then S both can, as of t , do both x at t' and y at t'' . More formally:

$$[O_t(x_{t'}) \ \& \ O_t(y_{t''])] \rightarrow C_t(x_{t'} \ \& \ y_{t''])$$

III. Portmore's Account of Moral Obligation

Recall the meta-criterion of rightness that I gave a couple lectures back:

MP S's performing x is *morally permissible* if and only if there is no available alternative that S has better requiring reason to perform and no worse reason, all things considered, to perform.

When we modify MP so as to accommodate the idea underlying FCAU, we get the following account of moral obligation:

MP* It is, as of t_n , morally permissible for S to perform x at t_m ($m > n$) if and only if there exists an FCA available to S at t_n in which S performs x at t_m for which there is no alternative FCA available to S at t_n in which S doesn't perform x at t_m that S has both better requiring reason and no worse reason, all things considered, to undertake.

AN ILLUSTRATION: Let's suppose that of all the FCAs available to me at t_0 , the one that I have most reason, all things considered, to undertake is one in which I work on my book at t_1 . And let's suppose that of all the FCAs available to me at t_0 , the one that I have most requiring reason to undertake is one in which I volunteer for Oxfam at t_1 . In this case, then, I'll have the moral option both to work on my book at t_1 and to volunteer for Oxfam at t_1 .

ANOTHER ILLUSTRATION: Here's a modified version of the example that I gave two lectures back. Let $MR(x)$ be a measure of how much moral reason the agent has to perform x , and let $ATCR(x)$ be a measure of how much reason, all things considered, the agent has to perform x :

Act x	Description of x	$MR(x)$	$ATCR(x)$
a ₁	spend Saturday volunteering for Oxfam	100	110
a ₂	spend Saturday volunteering for the ACLU	20	120
a ₃	spend Saturday working on my book	5	150
a ₄	spend Saturday watching TV ¹	0	5

Whereas MP entails that a₄ is morally wrong, MP* does not. Arguably, all the FCAs available to me at t_0 that I have most reason, all things considered, to undertake are ones in which I strike a certain balance between the time that I spend relaxing, the time that I spend working to help others, and the time that I spend working on my own personal projects. Thus, so long as I strike a reasonable balance between these three sorts of activities, it doesn't much matter which of the three I spend this Saturday doing. This means that when we look at FCAs as opposed to individual acts, things might look different than above. They might look like this:

FCA	Description of the FCA for Saturday	$MR(FCA_n)$	$ATCR(FCA_n)$
FCA ₁	spend Saturday volunteering for Oxfam	100	130
FCA ₂	spend Saturday volunteering for the ACLU	50	150
FCA ₃	spend Saturday working on my book	50	150
FCA ₄	spend Saturday watching TV	50	150

Imagine, then, that FCA₂, FCA₃, and FCA₄ are the same with respect to the time that I spend working on my book, volunteering for the ACLU, and watching TV. They differ only with respect to when I do these various activities. This is why both $MR(FCA_n)$ and $ATCR(FCA_n)$ are the same for each of the three instances. Thus MP* allows us to account for even more options than MP does.

Ultimately, then, I think that Commonsense Consequentialism should be formulated as follows:

CSC It is, as of t_n , morally permissible for S to perform x at t_m ($m > n$) if and only if there exists an FCA available to S at t_n in which S performs x at

¹ Watching TV on Saturday has the consequence of bringing me so pleasure, but that's about it. It does not have as a consequence that I will spend more of my time later working on my book or volunteering for the ACLU. This is why there is no moral reason and very little reason, all things considered, to watch TV this Saturday. In this respect, it's analogous to the act of Feldman's gathering his gardening tools. This has some negative utility because Feldman doesn't enjoy gathering his gardening tools, and it has no positive utility because Feldman's gardening is not a consequence of his gathering his gardening tools.

t_m for which there is no alternative FCA available to S at t_n in which S doesn't perform x at t_m whose outcome S has both better requiring reason and no worse reason, all things considered, to want to obtain.

So we look at the alternative FCAs available to S and what their outcomes will be. We then rank these outcomes with respect to both how much requiring reason S has to desire them and according to how much reason, all things considered, S has to desire them. And we judge that act is permissible if and only if it is deemed permissible by CSC.

IV. How the Portmore's Account of Moral Obligation Differs from Both Feldman's Account and Zimmerman's Account

A. Feldman's Account of Moral Obligation

"[A person] s morally ought, as of a time, t , to see to the occurrence of a state of affairs, p , iff p occurs in some world accessible to s at t , and it is not the case that $\sim p$ occurs in any accessible world as [intrinsically] good as (or better than) that one" (as quoted in Zimmerman 2006, 151).

Unlike Feldman's account, my account doesn't assume that moral permissibility of acts is a function solely of intrinsic value.

Unlike Feldman's account, my account doesn't have the odd implication that I am, as of t , morally obligated to see to the occurrence of any state of affairs, p , that is unalterable as of t . For instance, on Feldman's account, I am, as of now, morally obligated to see to the occurrence of the holocaust. And I am, as of now, morally obligated to see to the occurrence of the sun's rising tomorrow. Of course, these obligations are empty in that they don't require me to do any one thing over any other, for I will see to their occurrences no matter what I do. Nevertheless, it is an odd implication of Feldman's account and, it seems, an entirely unnecessary one.

Unlike Feldman's account, my account concerns whether or not an agent ought to perform some act rather than whether or not an agent ought to see to the occurrence of some state of affairs. This is important as it is not always clear what the relationship between the two is, as Zimmerman explains:

Suppose that Roger ought to raise his hand. Just what [state of affairs] is it that Roger ought to see to the occurrence of, his raising his hand or his hand's rising? In neither case is it clear that we get an equivalent claim; in one case too much seems to be said, in the other too little. To say that Roger ought to see to it that he raises his hand seems to say too much, in that it suggests that he ought to do something over and above just raising his hand. To say that

Roger ought to see to it that his hand rises seems to say too little, in that it appears compatible with his fulfilling his obligation by way of ensuring that his hand rises in such a way that it is not he but someone or something else that raises it. (As quoted in Zimmerman 2006, 166-167).

B. Zimmerman's Account of Moral Obligation

Zimmerman's account of moral obligation is unlike Feldman's account and like my account in the above three respects. Indeed, I've modeled my account off of Zimmerman's. Here's Zimmerman's account:

"s morally ought at t to do a at t' iff there is a world accessible to s at t in which s does a at t' and whose deontic value, relative to s at t , is greater than that of any world accessible to s at t in which s does not do a at t' " (Zimmerman 2006, 152).

Zimmerman intends "'deontic value' to function as a place-holder for whatever, in the end, should turn out to be the factor, or set of factors, that determines what s ought at t to do" (Zimmerman 2006, 152).

My account is just a specific version of Zimmerman's account where I've described what the factors are that determine what s ought at t to do. Also, I find it a bit more intuitive to talk about future courses of action available to an agent at a time rather than worlds accessible to an agent at a time, but I don't see any substantive difference between the two views.

IV. Actualism vs. Possibilism

All the accounts of moral obligation that we've considered so far have been versions of possibilism. Some favor actualism over possibilism. Below, we'll consider the difference between the two views and why we should favor possibilism over actualism.

ACTUALISM: S morally ought to do x if and only if S can do x and what *would* happen if S did x is better than what *would* happen if S did not do x . (Zimmerman 2006, 152)

POSSIBILISM: S morally ought to do x if and only if S can do x and what *could* happen if S did x is better than what *could* happen if S did not do x . (Zimmerman 2006, 153)

A. A Case that Illustrates How the Two Differ

To see how these two differ, consider the following case from Zimmerman:

Case 1:

I have been invited to attend a wedding. The bride-to-be is a former girlfriend of mine; it was she who did the dumping. Everyone, including me in my better moments, recognizes that she was quite right to end our relationship; we were not well suited to one another, and the prospects were bleak. Her present situation is very different; she and her fiancée sparkle in one another's company, spreading joy wherever they go. This irks me to no end, and I tend to behave badly whenever I see them together. I ought not to misbehave, of course, and I know this; I could easily do otherwise, but I do not. The wedding will be an opportunity for me to put this sort of boorishness behind me, to grow up and move on. The best thing for me to do would be to accept the invitation, show up on the day in question, and behave myself. The worst thing would be to show up and misbehave; better would be to decline and not show up at all. (Zimmerman 2006, 153)

Now suppose that what I would do if I were to accept this invitation is to show up and misbehave. I could accept, show up, and behave, but this is not what I would do if I were to accept the invitation. What should I do, then? Should I accept the invitation?

The possibilist says that I should accept, whereas the actualist says that I should decline.

TWO OTHER CASES TO CONSIDER: Feldman's case involving the doctor, the patient, and medicines A and B. Greenspan's case of the placement director.

B. Actualism's Commitments and Their Counterintuitive Implications

The actualist is committed to all of the following:

- (1) I can both accept the invitation and behave myself at the wedding. (This actually isn't something actualism implies, but it is a stipulation of the case and thus something that the actualist cannot deny.) [The possibilist accepts this.]
- (2) I cannot accept the invitation and behave myself at the wedding if I decline the invitation. (This actually isn't something actualism implies, but it is a fact that the actualist cannot deny.) [The possibilist accepts this.]
- (3) I ought both to accept the invitation and to behave myself at the wedding. [The possibilist accepts this.] This is so, for this is something I can do, and what would happen if I did is better than what would happen if I didn't.
- (4) It is not the case that I ought to accept the invitation. [The possibilist denies this this.] This is so, for what would happen if I did (namely, my misbehaving) is worse than what would happen if I didn't.

- (5) I ought to decline the invitation. [The possibilist denies this.] This is so, for would happen if I did is better than what would happen if I didn't.
- (6) I cannot accept the invitation and behave myself at the wedding without accepting the invitation. (This actually isn't something actualism implies, but it is a fact that the actualist cannot deny.) [The possibilist accepts this.]

And given these commitments, the actualist must deny the following plausible deontic principles:

- (a) *The Ought-Distributes-through-Conjunction Principle*: If S ought to do both x and y , then S both ought to do x and ought to do y . More formally:

$$O(x \ \& \ y) \rightarrow O(x) \ \& \ O(y)$$

To see that the actualist must deny this principle, consider the following. Let x = accept the invitation, and let y = behave myself at the wedding. (3) and (4) imply:

$$O(x \ \& \ y) \ \& \ \sim O(x)$$

- (b) *A Corollary of the Ought-Implies-Can Principle*: If S both ought to do x and S ought to do y , then S can do both x and y . More formally:

$$[O(x) \ \& \ O(y)] \rightarrow C(x \ \& \ y)$$

To see that the actualist must deny this principle, consider the following. Let x = accept the invitation and behave, and let y = decline the invitation. (3), (5), and (2) imply:

$$[O(x) \ \& \ O(y)] \ \& \ \sim C(x \ \& \ y)$$

- (c) *The Ought-Extends-to-Prerequisite-Acts Principle*: If S ought to do x but cannot do x without doing y , then S ought to do y . More formally:

$$[O(x) \ \& \ \sim C(x \ \& \ \sim y)] \rightarrow O(y).$$

To see that the actualist must deny this principle, consider the following. Let x = accept the invitation and behave, and let y = accept the invitation. (3), (6), and (4) implies:

$$[O(x) \ \& \ \sim C(x \ \& \ \sim y)] \ \& \ \sim O(y).$$

The possibilist can accept (a)-(c), for, unlike the actualist, the possibilist denies (4) and (5), which are the two claims that get the actualist in trouble.

C. Accounting for the Intuition that Seems to Support Actualism: Conditional Obligations, Wide-Scope Oughts, and Factual Detachment

The actualist might argue for her conclusion concerning Case 1 in the following way:

- i. I ought to decline the invitation *if* I am going to misbehave. (By intuition.)
- ii. I am going to misbehave. (By stipulation)
- iii. Therefore, I ought to decline the invitation. (By what's known as *factual detachment*: the detachment an unconditional obligation from both a fact, such as the one stated in (ii), and a conditional obligation, such as the one stated in (i).)²

Greenspan argues that we can accept both (i) and (ii) while denying (iii), for she thinks that (i) should be read as what's called a wide-scope ought (Broome 2004, 29). Let x = I am going to misbehave, and let y = I decline the invitation. Greenspan argues that (i) should be read as:

(i') $O(x \rightarrow y)$, or, equivalently, $O(\sim x \vee y)$ ³

Thus, we should not read (i) as:

(i'') $x \rightarrow O(y)$, or, equivalently, $\sim x \vee O(y)$

From (i'') and (ii), (iii) follows. But (iii) does not follow from (i') and (ii). Why think that (i') is the better reading? The answer is that there seems to be two ways that I could fulfill the obligation stated in (i) even on the assumption that I am going to misbehave. I could fulfill this obligation (to decline the invitation if I am going to misbehave) either by accepting and not misbehaving or by declining. As it happens, I ought to fulfill this obligation by doing the former as opposed to the latter. If, by contrast, (i'') was the better reading, then we should think that the only way for me to fulfill this obligation given that I am going to misbehave is to decline the invitation. But even if I am going to misbehave, I ought to accept the invitation and behave myself, as that would be best and that is something that I can do. We should not sanction wrongdoing, and we should not tailor our moral obligations to our moral failings.

² See Patricia S. Greenspan, "Conditional Oughts and Hypothetical Imperatives," *Journal of Philosophy* 72 (1975): 259-276.

³ Here 'O(x)' stands for 'I ought to act so as to ensure that x'.

There are other examples where the ought-operator seems to govern a conditional statement. Consider this one from Broome (2004, 29):

No one ought to believe that the world was made in less than a week; the evidence is strongly against it. Even if you believe that the world was made in six days, still it is not the case that you ought to believe it was made in less than a week. Nevertheless, you ought (to believe the world was made in less than a week, if you believe that it was made in six days). You can satisfy this requirement either by not believing the world was made in six days or by believing it was made in less than a week. As it happens, you ought to satisfy it the first way. You ought not to believe the world was made in six days, even if you do. (Broome 2004, 29)

Another example concerns mean-end reasoning. In the case of means-end reasons the ought seems to have wide scope. Thus the requirement of means-end reasoning reads as follows: "You ought (to intend to *M* if you intend to *E* and you believe that *M*-ing is a necessary means to your *E*-ing)" (Broome 2004, 29). So suppose that you intend to become a notorious killer and you believe that you must kill at least twenty people to become truly notorious. Does this mean that you ought to kill at least twenty people? Certainly, not!

If this right, then we can accept (i), which does seem intuitively plausible, without accepting (iii) and all of its absurd implications, which violate various plausible deontic principles. So it seems, then, that actualism has nothing going for it and a lot going against it. We should, therefore, accept possibilism, not actualism.