

Is Global Consequentialism More Expressive Than Act Consequentialism?

*Elliott Thornley**

Abstract: Act consequentialism states that an act is right iff the expected value of its outcome is at least as great as the expected value of any other act's outcome. Two objections to this view are as follows. The first is that act consequentialism cannot account for our normative ambivalence in cases where agents perform the right act out of bad motives. The second is that act consequentialism is silent on questions of character: questions like 'What are the right motives to have?' and 'What kind of person ought I be?'. These objections have been taken to motivate a move to global consequentialism, on which acts are not the only subjects of normative assessment. Motives and decision-procedures (amongst other things) are also judged right or wrong by direct reference to their consequences. In this paper, I argue that these objections fail to motivate the move from act to global consequentialism.

1. Introduction

Consequentialist theories begin with an axiology: a betterness ordering over outcomes. To this axiology, act consequentialists add the claim that acts are subject to normative assessment. An act is *right* iff the expected value of its outcome is at least as great as the expected value of any other act's outcome. An act is *wrong* otherwise. And, on act consequentialism, acts are the *only* subjects of normative assessment. Terms like 'right', 'wrong', 'obligatory', and 'permissible' apply to acts and acts alone.

Global consequentialists go further. On the purest form of the view, *everything* is normatively assessed by direct reference to the expected value of its outcome (Pettit and Smith 2000). Climates, eye colours, and constellations are either right or wrong, just as acts are. A more moderate global consequentialism demurs from these judgements, but nevertheless takes motives and decision-procedures (amongst other things) to be subjects of normative assessment (Feldman 1993; Louise 2006; Driver 2014; Greaves 2020). A set of motives, for example, is *right* for a person to have iff the expected value of having those

* Comments and questions welcome at elliott.thornley@philosophy.ox.ac.uk

motives is at least as great as the expected value of having any other set of motives. Having a set of motives is *wrong* otherwise.

Why extend normative assessment beyond acts? The main reason offered in the literature is that the resulting theory has more expressive power (Ord 2008, 31, 50–51). In particular, it’s been argued that global consequentialism, but not act consequentialism, can account for our *normative ambivalence* in cases where agents perform the right act out of bad motives (Louise 2006, 66, 84; Driver 2014, 168; Greaves 2020, 438). And it’s been argued that act consequentialism, but not global consequentialism, is silent on *questions of character*: questions like ‘What are the right motives to have?’ and ‘What kind of person ought I be?’ (Feldman 1993, 212; Louise 2006, 68; Driver 2014, 167–70; Greaves 2020, 436).

In this paper, I argue that these advantages are illusory. If certain empirical hypotheses about the persistence and causal efficacy of motives and decision-procedures are true, then *both* species of consequentialism can account for our normative ambivalence and offer responses to questions of character. If these hypotheses are false, then *neither* species can. So, whether the hypotheses are true or not, the *Normative Ambivalence* and *Character* objections give us no reason to prefer global consequentialism.

2. Normative Ambivalence and Character

Consider a case slightly adapted from Thomson (1976, 206):

Sarah – a doctor – kills a healthy person. She uses that person’s organs to save the lives of five patients.

Granted we fill out the details of this case in a certain way (Sarah knows that no one will find out, that the healthy person has no family or friends, etc.), act consequentialists will say that Sarah’s act is right. And since on their view normative assessment applies only to acts, they seem committed to the claim that nothing is wrong in this situation. But, as many point out, there’s something deeply amiss about the agents in these kinds of cases (Driver 2014; Greaves 2020). Normative ambivalence strikes us as a more appropriate response than whole-hearted approval.

Global consequentialists can account for our normative ambivalence. They too will judge that Sarah’s act is right. But her performing that act suggests that her motives are ruthless, and global consequentialists can say that these are the wrong motives. After all, motives are persistent, and ruthless motives are much less felicitous than compassionate ones in most situations. That makes plausible the claim that having compassionate motives has greater expected value, even taking into account the chance that compassion would have prevented Sarah from killing the unwilling organ donor. Global consequentialists might say something similar about Sarah’s decision-procedure. Her act suggests that her decision-

procedure is calculating; a spontaneous decision-procedure has greater expected value; and hence Sarah's act suggests that she has the wrong decision-procedure. These global consequentialist judgements – that Sarah's act is right, but her motives and decision-procedure are wrong – accord better with our intuitive appraisal of the case than the act consequentialist verdict.

Act consequentialism is also alleged to be incomplete. More specifically, it's said to be silent on *questions of character*: questions about the normative status of motives and decision-procedures. Since act consequentialism restricts normative assessment to acts, it cannot offer satisfactory responses to character questions like 'What are the right motives to have?' and 'What decision-procedure ought I have?'. Global consequentialism, on the other hand, takes these character questions in its stride. Agents ought to have those motives and that decision-procedure with the greatest expected value.

So much for the criticisms of act consequentialism. Now for two possible responses. The first is the evaluative strategy. Remember that all consequentialist theories begin with an axiology: a betterness ordering over outcomes. That allows act consequentialists to respond to each objection in evaluative terms. They can account for our normative ambivalence by claiming that, although Sarah's act is *right*, her motives and decision-procedure are *bad*. It would be better overall if she had compassionate motives and a spontaneous decision-procedure. Act consequentialists can also offer evaluative responses to the character questions. To the question, 'What are the right motives to have?', they can reply, 'Having compassionate motives has greatest expected value.'

This response may satisfy some. However, the character questions are *normative* so global consequentialists might well claim that only a normative answer will do. What's more, as the popularity of deontological theories attests, it's not obviously incoherent to suppose that what's best differs from what's right. Having identified the motives with the greatest expected value, there remains a further question whether these are the right motives.¹ Global consequentialists might say something similar with respect to normative ambivalence. Our worry about Sarah killing the donor is not just that something is *bad*; it's that something is *wrong*.

That brings us to the second possible response. Act consequentialists can try to *domesticate* global consequentialism (Railton 1984, 159–60; 1988, 403; Gruzalski 1986, 776). More precisely, they can replace the state-verbs in global consequentialist judgements with action-verbs, and hence bring correlates of these expressions into the act consequentialist fold. Whenever global consequentialists say, 'Sarah ought to *have* compassionate motives,' act consequentialists can say,

¹ As Chappell (2012) and McElwee (2020) point out, one can be a consequentialist about acts and yet claim that the normative status of non-acts depends on factors besides their expected value, like their intrinsic fittingness.

‘Sarah ought to *develop* compassionate motives.’ This manoeuvre yields a kind of normative ambivalence. By continuing to have ruthless motives, Sarah is failing to do what she ought. The domestication strategy also allows act consequentialists to give normative responses to the character questions. To the question, ‘What kind of person ought I be?’, they can reply, ‘You ought to *become* X kind of person.’

But, as Ord points out, there’s an issue here (2009, 23). Developing new motives and adopting new decision-procedures takes time. And this time-lag between acts and their fruits renders the act consequentialist domestication of global consequentialist judgements imperfect. We can illustrate the issue as follows. Granted we fill out the case in a certain way, global consequentialists will say that Sarah ought to *have* compassionate motives even as she’s preparing to kill the unwilling donor. Such motives might prevent Sarah from killing, but they have greatest expected value overall because a compassionate Sarah would perform better acts in later situations. Per the domestication strategy above, act consequentialists should swap the state-verb for an action-verb. But Sarah certainly ought not choose this time to *develop* compassionate motives. That’s because developing compassionate motives plausibly requires (a) compassionate acts and (b) time. So, if Sarah chose this time to develop compassionate motives, she would fail to perform the right act of killing the donor *and* fail to perform the right act in many later situations, in virtue of her (still) ruthless motives. The act consequentialist account of normative ambivalence thus evaporates. In killing the donor, Sarah does nothing wrong. Global consequentialism turns out to be harder to domesticate than we thought.

Nevertheless, I argue, the domestication strategy is a winning one for act consequentialists. Although time-lag scuppers the naïve version outlined above, the persistence and causal efficacy of motives and decision-procedures allows for a more sophisticated kind of domestication. In brief, act consequentialists can account for our normative ambivalence by noting that Sarah’s act is evidence that she has acted and will act wrongly at other times. And they can respond to the character questions by pointing out that some motives and decision-procedures are more conducive to right acts than others.

I develop these ideas below, using motives to illustrate. Decision-procedures can be treated in a similar manner.²

3. Empirical Hypotheses

Here are three empirical hypotheses:

² Plausibly, it’s easier to adopt new decision-procedures than new motives. If so, time-lag is less of a problem for decision-procedures. Any cases where time-lag remains an issue can be handled in the way I outline below. I thank an anonymous referee for pressing me to address this point.

Motives depend on acts

We develop compassionate (ruthless) motives by performing compassionate (ruthless) acts in ordinary situations.

Motives affect acts

The more compassionate (ruthless) a person's motives, the more likely they are to perform a compassionate (ruthless) act.

Extraordinary acts, extraordinary motives

A person will kill the unwilling organ donor iff they have extraordinarily ruthless motives.

And here are some assumptions to simplify the discussion. Each act is either compassionate or ruthless. Compassionate motives and ruthless motives are mutually exclusive, and each person starts out poised between the two. In ordinary situations, the right act by consequentialist lights is compassionate. In extraordinary situations like Sarah's, the right act is ruthless. Most situations are ordinary.

If the above three hypotheses are true, Sarah's act of killing allows us to draw inferences about her past and future acts. 'Extraordinary acts, extraordinary motives' lets us infer that Sarah has extraordinarily ruthless motives. 'Motives depend on acts' then lets us infer that Sarah has performed ruthless acts in ordinary situations, and so has acted wrongly in the past. 'Motives affect acts' lets us infer that Sarah will likely continue to perform ruthless acts in ordinary situations, and so will likely continue to act wrongly in the future. Act consequentialists can thus account for our normative ambivalence. Sarah is right to kill the donor, but whole-hearted approval seems inappropriate because her act is evidence that her past and future acts are wrong.

The truth of the above hypotheses would also allow act consequentialists to offer new responses to character questions like 'What motives ought I have?'. Given 'Motives depend on acts,' act consequentialists can say,

- (1) If you act as you ought in ordinary situations, you will have compassionate motives.

Given 'Motives affect acts,' act consequentialists can add,

- (2) If you have compassionate motives, you're more likely to act as you ought.

And, in cases where time-lag is not an issue, they can say,

- (3) You ought to develop compassionate motives.

Global consequentialists might object that these responses are still too roundabout. Not one is as simple and direct as the global consequentialist answer:

‘You ought to have compassionate motives.’ But this answer’s simplicity and directness strike me as dubious virtues. Were you to ask, ‘Why ought I have compassionate motives?’, a large part of the global consequentialist response would involve repeating the act consequentialist’s response (2) above.³ What’s more, if the global consequentialist answer is to be of practical interest, it must offer guidance (Gruzalski 1986, 775–76). And, as far as guidance goes, the act consequentialist’s response (3) matches the global consequentialist’s: one can abide by the instruction to *have* compassionate motives only by *developing* compassionate motives. For these reasons, I claim that the act consequentialist’s responses to the character questions are (at least) as good as the global consequentialist’s response.

Of course, the empirical hypotheses I give above are too strong. ‘Motives depend on acts’ and ‘Extraordinary acts’ are framed as certainties, but hypotheses about the connection between acts and motives can be at best probabilistic. With that in mind, we can weaken ‘Motives depend on acts’ as follows:

Motives depend on acts (probabilistic)

Performing compassionate (ruthless) acts in ordinary situations makes a person *more likely* to develop compassionate (ruthless) motives.

‘Motives affect acts’ is already probabilistic. I repeat it here for ease of reference:

Motives affect acts (probabilistic)

The more compassionate (ruthless) a person’s motives, the more likely they are to perform a compassionate (ruthless) act.

We can weaken ‘Extraordinary acts’ as follows:

Extraordinary acts, extraordinary motives (probabilistic)

A person with extraordinarily ruthless motives is *more likely* to kill the unwilling donor.

Now what can we infer from Sarah’s act of killing? The probabilistic version of ‘Extraordinary acts’, plus Bayes’ theorem, implies that we should increase our credence that Sarah has extraordinarily ruthless motives.⁴ Given the probabilistic

³ Only ‘a large part’ because some of the consequences of motives and decision-procedures are not mediated by acts. I address this point below.

⁴ Let A = Sarah killed the donor. Let B = Sarah has extraordinarily ruthless motives. ‘Extraordinary acts’ states that $P(A) < P(A|B)$. Bayes’ theorem states that $P(A|B) = P(B|A) \cdot P(A) / P(B)$. Substituting $P(A|B)$ gives $P(A) < P(B|A) \cdot P(A) / P(B)$. Rearranging and

version of ‘Motives depend on acts,’ that in turn implies that we should increase our credence that Sarah has acted wrongly at least m times in the past, for any positive integer m .⁵ Given the probabilistic version of ‘Motives affect acts,’ we should increase our credence that Sarah will act wrongly at least n times in the future, for any positive integer n .⁶ The upshot is that even these probabilistic hypotheses yield an act consequentialist account of our normative ambivalence. Sarah’s act of killing is right, but it’s evidence that her past and future acts are wrong.

The probabilistic hypotheses also permit act consequentialist responses to the character questions. To the question, ‘What motives ought I have?’, act consequentialists can respond,

(1*) If you act as you ought in ordinary situations, you will likely have compassionate motives.

(2*) If you have compassionate motives, you’re more likely to act as you ought.

And in cases where time-lag is not an issue,

(3*) You ought to develop compassionate motives.

Now suppose that even the probabilistic version of ‘Motives depend on acts’ is not true of Sarah: performing compassionate (ruthless) acts makes Sarah no more likely to form compassionate (ruthless) motives. In that case, act consequentialists can still account for our normative ambivalence. The probabilistic versions of ‘Extraordinary acts’ and ‘Motives affect acts’ still imply that, after learning Sarah killed the donor, we should increase our credence that Sarah will perform at least n wrong acts in the future. And act consequentialists can still offer (2*) and (3*) as responses to the character questions.

Suppose instead that even the probabilistic version of ‘Motives affect acts’ is not true of Sarah: Sarah’s probability of performing a compassionate versus ruthless act in an ordinary situation is the same no matter what her motives. In this case too, act consequentialists can account for our normative ambivalence. The probabilistic versions of ‘Extraordinary acts’ and ‘Motives depend on acts’

cancelling $P(A)$ gives $P(B) < P(B|A)$. In other words, learning that Sarah killed the donor should increase our credence that she has extraordinarily ruthless motives.

⁵ Let B = Sarah has extraordinarily ruthless motives. Let C = Sarah has performed at least m wrong acts. ‘Motives depend on acts’ states that $P(B) < P(B|C)$. By Bayes’ theorem, $P(C) < P(C|B)$. In other words, increasing our credence that Sarah has extraordinarily ruthless motives should increase our credence that Sarah has performed at least m wrong acts.

⁶ Let B = Sarah has extraordinarily ruthless motives. Let D = Sarah will perform at least n wrong acts in the future. ‘Motives affect acts’ implies that $P(D) < P(D|B)$. In other words, increasing our credence that Sarah has extraordinarily ruthless motives should increase our credence that Sarah will perform at least n wrong acts in the future.

still imply that we should increase our credence that Sarah has performed at least m wrong acts in the past. And act consequentialists can still offer (1*) and (3*) as responses to the character questions.

Now suppose that neither ‘Motives depend on acts’ nor ‘Motives affect acts’ is true of Sarah: in ordinary situations, Sarah’s acts have no effect on her motives and her motives have no effect on her acts. In that case, we cannot infer anything about Sarah’s past or future acts in ordinary situations from her act of killing, so the act consequentialists’ account of our normative ambivalence evaporates. The same goes for their three responses to the character questions.

But note two points. The first is that this implication might be just what we want. If Sarah’s motives are entirely inert in ordinary situations, then whole-hearted approval of Sarah’s act does not seem inappropriate, and the character questions seem to lack any sensible answer. The second is that, in any case, global consequentialists are in a similar predicament. Set aside for a moment any consequences of motives not mediated by acts.⁷ If Sarah’s motives are inert in ordinary situations, then the expected value of her having compassionate motives is no greater than the expected value of her having ruthless motives. That means that the global consequentialists’ account of our normative ambivalence and their answers to the character questions also evaporate.

Now suppose instead that the probabilistic version of ‘Extraordinary motives’ is not true of Sarah: having extraordinarily ruthless motives makes Sarah no more likely to kill the donor. In that case, we cannot infer anything about Sarah’s motives from her act of killing. That in turn implies that we cannot infer anything about her past and future acts, so the act consequentialist’s account of our normative ambivalence evaporates. But, again, so does the global consequentialist’s: having ruthless motives might be wrong, but Sarah’s act is no evidence that she has such motives. As before, act consequentialism matches global consequentialism for expressive power.

I can now discuss a matter twice-delayed. Thus far, I have assumed that all the consequences of motives are mediated by acts, but that is not true (Pettit and Smith 2000, 122; Louise 2006, 67; Chappell 2012, 697–98). Even if having compassionate motives would not affect Sarah’s acts in ordinary situations, having them might have greater expected value than having ruthless motives in virtue of the joy that they inspire in Sarah. If this joy is great enough, global consequentialists will judge that Sarah ought to have compassionate motives.

My response is that the dialectic is the same here as elsewhere. In cases where time-lag is not an issue, act consequentialists can substitute an action-verb in place of the global consequentialist’s state-verb: Sarah ought to *develop*

⁷ Again, I address this point below.

compassionate motives. In continuing to have ruthless motives, Sarah is failing to do as she ought.

Now consider a case where time-lag is an issue. Sarah has killed the donor as before, but this time our empirical hypotheses are different. Instead of ‘Motives depend on acts,’ we have:

Changing motives takes time

We can change our motives only gradually over time.

Instead of ‘Motives affect acts,’ we have:

Motives affect joy

The more compassionate our motives, the more we enjoy each day.

Granted that the joy arising from compassion is great enough, global consequentialists can say that Sarah ought to *have* compassionate motives. But act consequentialists cannot say that Sarah ought to *develop* compassionate motives. Given ‘Extraordinary acts,’ that would prevent her from killing the donor, and given ‘Changing motives takes time,’ she would still have ruthless motives for a period afterward. And this time act consequentialists cannot say that Sarah’s act is evidence that she has acted and will act ruthlessly in ordinary situations because, *ex hypothesi*, her acts are unaffected by her motives. So, we might think, global consequentialists can account for our normative ambivalence and answer the character questions in this case, while act consequentialists cannot.

But note that, as before, Sarah’s act is evidence that she has extraordinarily ruthless motives. And given ‘Changing motives takes time,’ that in turn is evidence that Sarah has acted wrongly in the past by failing to develop compassionate motives. And insofar as a person’s past acts are a guide to their future acts, it is evidence that she will continue to act wrongly in this way. Once again, we have an act consequentialist account of our normative ambivalence. The same point allows act consequentialists to respond to the character questions. They can say, ‘If you act as you ought, you will develop compassionate motives.’

Now consider a final case.⁸ Sarah exists for just a moment. Her only act is right, but her motives are ruthless, and compassionate motives would give her greater joy. In this case, global consequentialism implies that we should be normatively ambivalent: Sarah’s act is right, but her motives are wrong. Act consequentialism, meanwhile, implies that we shouldn’t be normatively ambivalent: Sarah never acts wrongly. One might take this case as reason to prefer global consequentialism.

⁸ I thank an anonymous referee for bringing this case to my attention.

I claim, however, that this case is better understood as a point in favour of act consequentialism. That's because normative ambivalence seems inappropriate: Sarah's motives are certainly unfortunate, but nothing about her seems *wrong*. What's more, act consequentialism gives us the reason why: Sarah *does* nothing wrong. She doesn't choose her ruthless motives, and she's so short-lived that she can do nothing to change them.

4. Conclusion

Global consequentialists argue that their view outstrips act consequentialism with respect to expressive power. In particular, they claim that global consequentialism – but not act consequentialism – can account for our normative ambivalence in cases where agents perform the right act out of bad motives, and that act consequentialism – but not global consequentialism – is silent on questions of character. In this paper, I have argued that both claims are false. If certain empirical hypotheses about the persistence and causal efficacy of motives and decision-procedures are true, act consequentialists can account for our normative ambivalence by noting that the agent's act is evidence that they have acted and will act wrongly at other times. And they can respond to the character questions by citing the connections between certain motives, decision-procedures, and right acts. If these empirical hypotheses are not true, then the global consequentialist account of our normative ambivalence and answers to the character questions also evaporate. So, whether the hypotheses are true or not, the Normative Ambivalence and Character objections give us no reason to prefer global consequentialism.⁹

5. References

- Chappell, Richard Yetter. 2012. 'Fittingness: The Sole Normative Primitive'. *The Philosophical Quarterly* 62 (249): 684–704.
- Driver, Julia. 2014. 'Global Utilitarianism'. In *The Cambridge Companion to Utilitarianism*, edited by Ben Eggleston and Dale E. Miller, 166–76. Cambridge Companions to Philosophy. Cambridge: Cambridge University Press.
- Feldman, Fred. 1993. 'On the Consistency of Act- and Motive-Utilitarianism: A Reply to Robert Adams'. *Philosophical Studies* 70 (2): 201–12.

⁹ I thank Hilary Greaves, Tomi Francis, Alice van't Hoff, Kacper Kowalczyk, Aidan Penn, Nikhil Venkatesh and three anonymous referees for *Analysis* for helpful comments and discussion.

- Greaves, Hilary. 2020. 'Global Consequentialism'. In *The Oxford Handbook of Consequentialism*, edited by Douglas W. Portmore, 423–40. Oxford: Oxford University Press.
- Gruzalski, Bart. 1986. 'Parfit's Impact on Utilitarianism'. *Ethics* 96 (4): 760–83.
- Louise, Jennie. 2006. 'Right Motive, Wrong Action: Direct Consequentialism and Evaluative Conflict'. *Ethical Theory and Moral Practice* 9 (1): 65–85.
- McElwee, Brian. 2020. 'The Ambitions of Consequentialism'. *Journal of Ethics and Social Philosophy* 17 (2): 198–218.
- Ord, Toby. 2008. 'How to Be a Consequentialist about Everything'. Unpublished manuscript.
- . 2009. 'Beyond Action: Applying Consequentialism to Decision Making and Motivation'. DPhil Thesis, University of Oxford.
- Pettit, Philip, and Michael Smith. 2000. 'Global Consequentialism'. In *Morality, Rules, and Consequences: A Critical Reader*, edited by Brad Hooker, Elinor Mason, and Dale E. Miller, 121–33. Edinburgh: Edinburgh University Press.
- Railton, Peter. 1984. 'Alienation, Consequentialism, and the Demands of Morality'. *Philosophy and Public Affairs* 13 (2): 134–71.
- . 1988. 'How Thinking About Character and Utilitarianism Might Lead to Rethinking the Character of Utilitarianism'. *Midwest Studies in Philosophy* 13 (1): 398–416.
- Thomson, Judith Jarvis. 1976. 'Killing, Letting Die, and The Trolley Problem'. *The Monist* 59 (2): 204–17.