

# The Paralysis Argument

Ryan Doody

April 5, 2022

## Non-consequentialism

MacAskill & Mogensen argue that, given plausible assumptions about the long-run effects of everyday actions, **non-consequentialism** implies—absurdly—that we should do as little as possible in our lives. They call this the *Paralysis Argument*.

The argument proceeds by highlighting two standard assumptions of non-consequentialism:

- *Doing vs. Allowing*. The moral reasons against *doing harm* are stronger than the moral reasons against *allowing harm*.
- *Harm vs. Benefit*. The moral reasons against *doing harm* are stronger than the moral reasons in favor of *doing good* (even when the harms and benefits are of equal magnitude).

They also argue that following principle is false:

- *Inverse Doctrine of Doing and Allowing (IDDA)*. The moral reasons in favor of *doing good* are stronger than the moral reasons in favor of *allowing good* to happen.

Instead, they think, all else equal, your reasons in favor of benefiting others are not stronger than your reasons in favor of allowing people to be benefitted.

## Uncertainty

These are claims about *objective* moral reasons. But, because we're interested in cases in which we don't know exactly what the long-run effects of our actions will be, we need to look at whether these claims hold about *subjective* moral reasons.

**The Decision-Theoretic Approach.** The agent ranks all possible outcomes that might be realized by her actions in terms of the extent to which they would be favored by her objective moral reasons. (Call this their *choice-worthiness*.) The action that has the most *subjective* reason to be performed is the one that maximizes expected choice-worthiness.

The long-run effects of our everyday actions (e.g., going to the supermarket) are *identity-affecting*: they affect which persons belong to the total population of everyone who ever exists.

Non-consequentialists standardly endorse a constraint against *doing harm* (which is treated as importantly different from *allowing harm* and from *benefiting*).

They call this the *Doctrine of Doing vs Allowing (DDA)*.

They call this the *Harm Benefit Asymmetry (HBA)*.

They motivate this with the *Medicine* example: If you use the vial of medicine, you can save 1 life; if you allow others to use it, they can save 2 lives. Should you use it so that you can *do good*, rather than merely *allow* it?

*Objective moral reasons* are grounded in the way things actually are, not the agent's beliefs.

*Subjective moral reasons* are grounded in how the agent takes things to be, their beliefs, or evidence, or credences.

If  $\phi$  is as likely to result in outcome  $O$  as  $\psi$  is, and you have stronger moral reasons against  $\phi \rightsquigarrow O$  than against  $\psi \rightsquigarrow O$ , you thereby have stronger *pro tanto* subjective moral reason against  $\phi$  than against  $\psi$ . (These subjective reasons can be outweighed by others, at least until we've considered all possible  $O$ s.)

## The Paralysis Argument

Consider two actions ( $D$  and  $A$ ), and a morally significant outcome ( $O$ ) that might be one of their unforeseen consequences. Suppose  $O$  is harmful to somebody.

Let  $D$  be such that, if  $O$  results from  $D$ , you've *harmed* someone.

Let  $A$  be such that, if  $O$  results from  $A$ , you've *allowed* harm to happen.

You, then, have stronger objective moral reasons against performing  $D$  than  $A$  conditional on  $O$  being the result.

Suppose, instead, that  $O$  is beneficial. Because you don't have stronger moral reason to *do* good than to *allow* it, you don't have stronger objective moral reasons to favor  $D$  over  $A$  conditional on  $O$  being the result.

You have no more reason to think that  $O$  will result from  $D$  than from  $A$ , so you should regard the possibilities as equally likely. As we've seen,

If  $O$  is harmful, you have subjective reason to prefer  $A$ .

If  $O$  is beneficial, you don't have subjective reason to prefer  $D$ .

Because this holds for all  $O$ , you have greater subjective reason to favor  $A$  over  $D$ —that is, “to ensure, so far as possible, that the indirect consequences of your behavior are things you allow to happen, and not things you make happen.” (p. 6).

## Objections

1. *The Unforeseeable is Irrelevant.* The effects of your actions which are unforeseeable (even to the ideally conscientious agent) are morally irrelevant, and so they can be ignored.
2. *The Agency of Others.* If the causal chain from your action to a harmful outcome passes through the voluntary actions of other agents, your action isn't an act of *doing* harm.
3. *The Absolution of Inscrutability.* If the causal chain from your action to the harmful outcome is convoluted and inscrutable, then your action isn't an act of *doing* harm.
4. *Justifiable to All ex ante.* You don't have stronger moral reason against doing (rather than allowing) harm when that harm is *ex ante* justifiable to all.
5. *Undermined by Non-Identity.* Because of the Non-Identity Problem, the unforeseen consequences of your everyday actions are unlikely to *harm* anyone.
6. *Not Acting vs. Allowing.* Failing to act doesn't guarantee that the outcomes that will result are ones that you merely *allowed* to occur.

$D$ : Driving to the supermarket.

$A$ : Sitting motionless at home.

$O$ : The unpredictable premature death of someone many years from now.

This follows from  $DDA$ .

This follows from  $IDDA$ .

This invokes the *Principle of Indifference* (see our discussion of *Cluelessness*.)

They motivate the argument with a thought experiment:

*The Dice of Fortuna.* There is a sealed box containing dice of unknown nature and provenance. If you roll the dice and the result is *below* the average value of the sum of the numbered faces on the dice, a life will be *saved*. If the result is *above* the average, someone will be *killed*. You are offered \$1 to roll the dice. Should you?

Does this mean we should ignore outcomes with very small probabilities? (See: *The Button*)

Is that right? (See: *The Arms Trader*)

Is that right? (See: *The Mystery Box*)

M&M worry that many deontologists will reject the implications of this view (See: *The AI Ambulance*)

This response is only available to those who accept the Counterfactual Comparative Account of harming. Also, the harms of one's everyday actions can occur to already-existing people.

We need to think more about how to make the doing vs. allowing distinction.