

Absolutism, Utilitarianism and Agent-Relative Constraints

Mark T. Nelson

ABSTRACT: Absolutism—the idea that some kinds of acts are absolutely wrong and must never be done—plays an important role in medical ethics. Nicholas Denyer has defended it from some influential consequentialist critics who have alleged that absolutism is committed to “agent-relative constraints” and therefore intolerably complex and messy. Denyer ingeniously argues that, if there are problems with agent-relative constraints, then they are problems for consequentialism, since it contains agent-relative constraints, too. I show that, despite its ingenuity, Denyer’s argument does not succeed. The defense of absolutism must move to other grounds.

INTRODUCTION

WHEN THEN VICE-PRESIDENT George H. W. Bush was asked why he didn’t criticize President Reagan’s Iran-Contra activities, he is supposed to have replied, “Because you don’t tackle your own quarterback!” The implied principle seems sensible enough: it is wrong-headed to criticize people on your own side, pulling in the same direction as you.¹ Yet I intend to do just that: I am sympathetic to ethical absolutism, but I shall criticize an argument by one of its prominent defenders.²

The argument in question is from Nicholas Denyer’s intriguing essay, “Is Anything Absolutely Wrong?”³ In it, Denyer defends absolutism, the idea that some kinds of acts, such as infanticide, are absolutely wrong, meaning that one must never do them, regardless of the consequences.⁴ That is my rough definition;

¹We may wish to make an exception for cases where a confused quarterback starts to run the ball toward the wrong end-zone.

²Other defenders include Leon Kass, G. E. M. Anscombe, David Oderberg, David Solomon.

³Nicholas Denyer, “Is Anything Absolutely Wrong?” in *Human Lives: Critical Essays on Consequentialist Bioethics*, eds. David Oderberg and Jacqueline A. Laing (London: Macmillan, 1997), 39–57.

⁴A typical statement of this idea is by G. E. M. Anscombe: “But if someone really thinks, in advance, that it is open to question whether such an action as procuring the judicial execution of the innocent should be quite excluded from consideration—I do not want to argue with him; he shows a corrupt mind.” G. E. M. Anscombe, “Modern Moral Philosophy,” *Philosophy* 33 (1958): 1–19, at 16–7. An influential presentation

Denyer's is more precise. Observing that human actions can be described in many ways, he defines absolutism as:

the idea that certain descriptions (e.g., “dismembering a baby”) are such that any action satisfying any of those descriptions is for that reason wrong, whatever other descriptions (e.g., “saving a life”) it may also satisfy. This idea is called absolutism. It was once the conventional wisdom.⁵

It may no longer be the conventional wisdom, but it is still a staple of much philosophical thought about medical ethics, especially when combined with the doctrine of double effect and associated distinctions between causation, foresight and intention. It is absolutism in this sense to insist (as in Denyer's example) that a doctor must not dismember a baby, even when that is the only way to save the life of its mother; or that a doctor must never kill patients, even terminal ones suffering from chronic pain, who have requested it.⁶ Absolutism is not just practically important, it is theoretically important, too: some philosophers see it as a boundary between consequentialist and non-consequentialist ethical theories.⁷

Staple or not, absolutism has come in for withering criticism, especially from utilitarians and other consequentialists. Much of their criticism has been aimed at the outer layer of associated doctrines and distinctions, but some has been aimed at the core principle itself.⁸ According to Denyer:

Even if absolutists are conceded all the distinctions they wish to draw concerning causation and intention, there still remains another objection. How should you respond if someone threatens to dismember a dozen babies unless you dismember one? Suppose that dismembering babies is absolutely wrong. In that case, it would be wrong for you to dismember the one baby even if the threat were entirely credible; you must refuse to breach the absolute prohibition, even if your refusal means

of the idea specifically in medical ethics is Leon Kass, “Neither for Love nor Money: Why Doctors Must Not Kill,” *The Public Interest* 94 (1989): 25–46.

⁵Denyer, 39.

⁶Some consequentialist theories allow for principles that mimic absolutism up to a point, but that is another story. Also, it goes without saying that absolutism is a generic formal position involving absolute prohibitions. In itself, it is not committed to any particular prohibitions, such as those mentioned, against dismembering babies.

⁷See, e.g., Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986), 165. Not all non-consequentialists see it this way, however. For a dissenting view, see David McNaughton and Piers Rawling, “Contours of the Practical Landscape,” in *Thinking about Reasons: Themes from the Philosophy of Jonathan Dancy*, eds. David Bakhurst, Margaret Little, and Brad Hooker (Oxford: Oxford University Press, 2013), 240–64, esp. sections IV–VI. It should also be noted that I am taking “consequentialism” in its contemporary, narrower sense, and not in Anscombe's original, broader sense. Denyer himself calls attention to these two senses at Denyer, 39.

⁸For a lucid introduction to some recent criticisms of the doctrine of double effect and associated ideas concerning the moral relevance of intentions, see Helen Frowe, *The Ethics of War and Peace* (New York: Routledge, 2011), especially chapters 1 and 7.

in the present circumstances that the prohibition will be breached a dozen times. A technical term has been devised for this feature of absolutism: an absolutist ethic is said to contain “agent-relative constraints.” And the fact that absolutism contains agent-relative constraints is sometimes taken as a difficulty for it.⁹

A PROBLEM WITH AGENT-RELATIVE CONSTRAINTS

To see why a commitment to agent–relative constraints (hereafter ARCs) might be taken as a difficulty, we must consider them more closely. As the name suggests, an ARC is a *constraint* on the action of the agent, because it entails that, for some action kind, Φ , *agents are not to* Φ , even if by Φ -ing they could minimize the Φ -ing of others. This constraint is *agent-relative* in that it has the *agent* think, “I mustn’t Φ ,” rather than, “ Φ -ing is absolutely wrong, so I should try to minimize everyone’s Φ -ing overall.”¹⁰ Such constraints can be expressed in terms of wrongness, as above, or in terms of reasons for action. Thomas Nagel, e.g., describes constraining agent-relative *reasons* as “. . . not neutral reasons for everyone to bring it about that no one is maltreated, but relative reasons for each individual not to maltreat others himself, in his dealings with others.”¹¹

What is wrong with ARCs, thus described? According to Samuel Scheffler “. . . if the violation of agent-relative constraints is morally so objectionable, it seems extremely odd, on the surface at least, for morality to tell us that we must not act in such a way as to minimize their occurrence.”¹² There are different ways of spelling out this oddness, some more impressive than others. One unimpressive way to do it is to reason thus: “If Φ -ing is bad, then we have a moral obligation to minimize Φ -ing, regardless of the consequences. But to recognize and act on an ARC is precisely *not* to do this, therefore, to recognize and act on an ARC is to violate one’s moral obligation to minimize the bad, regardless of the consequences.” Of course, this just begs the question against deontology and against ARCs, by presupposing consequentialism.

A less obviously question-begging difficulty with ARCs is that they complicate our ethical theory. They complicate our ethical theory because they complicate our account of the relation between practical reasoning and value generally. One sort of practical reasoning about value which is supposedly familiar and understandable goes like this:

For some thing, property, or state of affairs, G :

1. I value G .

⁹Denyer, 39–57, at 48–49.

¹⁰The *locus classicus* of such cases is Bernard Williams’s discussion of integrity in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1963), 115f.

¹¹Nagel, 165.

¹²Samuel Scheffler, “Introduction,” in *Consequentialism and its Critics*, ed. Scheffler, (Oxford: Oxford University Press, 1988), 9, quoted in Denyer, 49.

Therefore,

2. I shall maximize G (where “maximize” means “produce as much G as possible” or “bring about a world such that no other possible world has more G in it”).

or

3. I disvalue G .

Therefore,

4. I shall minimize G (where “minimize” means “produce as little G as possible” or “bring about a world such that no other possible world has less G in it”).

For example, if I value my own pleasure (and nothing else), then, given instrumental rationality, there is only one appropriate response on my part: I should try to produce as much pleasure for myself as I can. If I disvalue pain (and nothing else), there is only one sort of appropriate response to that: I should try to produce as little pain for myself as I can. Whatever problems beset egoistic hedonism, theoretical complexity is not one of them, as it makes the relation between practical reason and value simple and clear.

The same cannot be said for accounts that admit ARCs, because they explicitly reject the idea that right actions always maximize value or minimize disvalue. One could, I suppose, simply detach practical reasoning about right action from considerations of value and disvalue, and hold that reasoning about good and bad is one thing, reasoning about what to do is quite another, and never the twain shall meet. This, however, flies in the face of the obvious fact that they sometimes *do* meet, that sometimes our practical reasoning is based on considerations of value and disvalue. Moreover, it cannot but diminish the degree of coherence and integration of our resultant ethical theory, as it replaces one unified account with two unrelated accounts.

Suppose, on the other hand, that we try for a unified account of practical reasoning and value. This will require us to express ARCs in terms of some new concepts, such that, sometimes, for some G s, the appropriate response is to try to maximize them, but that, at other times, for some *other* G , maximizing is not the appropriate response. That is, suppose we sometimes reason as follows:

1.' I value G .

Therefore,

2.' I shall *honor* G (where “honor” means something different from “maximize G ” or even “produce more G ”, and sometimes requires us *not* to do certain things even when these would lead to a world with more G in it).

or

3. I disvalue *G*.

Therefore,

4. I shall *oppose* *G* (where “oppose” means something different from “minimize *G*”, and requires us *not* to do certain things even when this leads to a world with more *G* in it).

On this account, if I value (or disvalue) *G*, the appropriate response is not necessarily the one suggested by instrumental rationality, that of producing more (or less) of it. After all, *G* may not be the sort of good or bad that can be “produced” in a way that permits of maximization or minimization. Or even if it can be, perhaps it shouldn’t be. Either way, this sort of practical reasoning about *G* dictates that some other response is appropriate. Recognizing the value of human life, e.g., does not mean producing as much of it as possible, but rather honoring it or protecting it, which often means *doing* some things (such as feeding children), and, equally, *not* doing other things (such as killing them).¹³

Anyway, some such conceptual devices are necessary if we are to unite our practical reasoning about action with our valuing of innocent lives (or our disvaluing of murder), but, whichever device we choose, we will be introducing a new category of value, a new kind of reasoning and a new, irreducibly different, pattern of appropriate response to value and disvalue. This is a problem, or at least a cost, if we conceive of our ethical theories in explanationist terms, as attempts to give unified explanatory theories, and evaluate them according to the standard criteria for good explanations.¹⁴ The standard criteria for good explanations always include something like “simplicity” and “internal coherence,” so any account that includes ARCs must inevitably fare worse on one or the other of these criteria than comparable theories that do not include ARCs. Either they will not unite their account of value with their account of practical reason, in which case they are less coherent; or they will unite their accounts, but at the cost of introducing new and complicating concepts and categories. In this case, says Philip Pettit, “the non-consequentialist [who accepts ARCs] is committed to a theory which is seriously defective in regard to the methodological virtue of simplicity.”¹⁵

¹³Philip Pettit makes exactly this argument, but in terms of “promoting” versus “honoring,” in “Consequentialism” in *Companion to Ethics*, ed. Peter Singer (Oxford: Blackwell, 1991), 230–40.

¹⁴For a lucid, recent presentation and defense of explanationism, see Ted Poston, “Explanationist Plasticity and the Problem of the Criterion,” *Philosophical Papers* 40 (2011): 395–419.

¹⁵Pettit, “Consequentialism”, 230–40, at 237. Pettit makes the criticism explicitly; Shelly Kagan makes it implicitly, when he complains that philosophers such as Nagel and Scheffler invoke some notion of agent-relativity without explaining it. For Kagan, “explaining” must be something more than merely “arguing for the existence of,” because Kagan objects to an argument of Nagel’s (purporting to show that preferences generate agent-relative reasons) saying, “Even if sound, I do not think such a *via negativa* will be especially illuminating. It may help establish the *existence* of agent-relative reasons, but it will not *explain* them.” I surmise that he means something like embedding the idea of agent-relativity in an overall theory

DENYER'S RESPONSE TO THE PROBLEM

How should absolutists respond to such a criticism? Since they can hardly deny that ARCs do entail this greater degree of complexity, one perfectly sensible response would be to point out that the criterion of theoretical simplicity is best understood as requiring, not the avoidance of all complexity, but the avoidance of *unnecessary* complexity. Moreover, simplicity is just one theoretical virtue among others, and there is no reason to privilege it over, say, conservativeness or explanatory comprehensiveness. Any ethical theory that does not contain ARCs either will be forced to jettison some important first-order moral judgments, or will retain those first-order judgments but leave them unexplained. That is, absolutists can just insist that a complex moral theory is needed to map a complex moral reality, and that, as W. D. Ross observes, “Loyalty to the facts is worth more than a symmetrical architectonic or a hastily reached simplicity.”¹⁶

Curiously, Denyer does not do this. He argues instead that, if there are problems with ARCs, then they are problems for consequentialism, since it contains ARCs, too.¹⁷ He frames his argument in terms of act (or direct) utilitarianism, the consequentialist theory par excellence, but his point is meant to apply to all direct versions of consequentialism as well. It is an ingenious argument, so worth quoting in full:

Absolutism is not unique in containing agent-relative constraints. Utilitarianism does too. This is because of what we might call the Dunkirk effect: people are normally rather lackadaisical about doing their best; but they show more zeal when desperate. Suppose I spend next weekend maximizing utility. This will, we may presume, mean leaving my close neighbours at large. They will potter round their gardens, each producing, let us say, 10 kilobentham of utility; and they will do this even though they could so arrange their affairs as to produce 11 kilobentham apiece instead. In other words, if I myself maximize utility and leave them at large, the consequence will be a dozen failures on their parts to maximize utility. Suppose now that instead of maximizing utility myself, I kidnap, bind and gag my dozen neighbours. In such a condition none of them will produce anything like 10 kilobentham; for it will be all they can do to keep breathing. But—and this is the important point—the dozen will under those circumstances do the best they can.

that meets the various criteria for explanatory goodness to a sufficiently high degree. See Shelly Kagan, “Does Consequentialism Demand Too Much? Recent Work on the Limits of Obligation,” *Philosophy and Public Affairs* 13 (1984): 239–54, at 246–7 (latter emphases mine). He complains similarly about the lack of explanation for agent-relative reasons elsewhere on pp. 246, 248, 249, 253, and 254. This is not the only criticism Pettit levels against ARCs; elsewhere he argues that non-consequentialism’s commitment to agent-relativity can be sustained only by denying the requirement of universalisability or by construing moral conclusions as agent relative as well. See Philip Pettit “Non-consequentialism and Universalizability,” *Philosophical Quarterly* 50, (2000): 175–90.

¹⁶W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930), 23.

¹⁷Perhaps it is not so curious. “Theoretical virtue” arguments are comparative, and the comparisons are inevitably person-relative. E.g., theories are not just “conservative,” full stop. They are conservative to the extent that they conserve some particular beliefs; to persons who hold those beliefs, the fact that a theory conserves them constitutes a virtue; to persons who do not hold those beliefs, it does not. This person-relativity means that theoretical virtue arguments are less likely to be dialectically decisive between radically divergent views.

For they will shun the less than maximally utile option of holding their breath; they will instead take the option of continuing to breathe, and none of them will have any third option more utile than that. Thus, if I kidnap my neighbours, there will be only one failure to maximize utility: my own. As a dutiful utilitarian, how do I spend my weekend? Do I maximize utility, and thus leave my neighbours free? Or do I kidnap the neighbours and thus minimize failures to maximize utility? Obviously, a dutiful utilitarian chooses the former. In other words, utilitarians believe in an agent-relative constraint: they believe it wrong to fail to maximize utility, even though by committing one such wrong oneself, one can prevent a dozen such wrongs committed by others, and even though failing to maximize utility is a wrong so horrendous that, when faced with a choice between it and murder, one should commit the murder. Hence if there is something objectionable about agent-relative constraints, the objection applies more widely than just to absolutism.¹⁸

DENYER'S RESPONSE ANALYZED AND CRITICIZED

Ingenious though this argument is, I do not believe it succeeds. Let us consider a simplified version of Denyer's scenario, involving only him and two neighbors, Patricia and Quentin.¹⁹ And let us focus, with Denyer, on the choice between two courses of action on a Saturday. Course of action *A* involves him walking to the nearby Oxfam office, and spending the day there, efficaciously volunteering. And let us grant that this alone, out of the range of actions open to him, would maximize utility. (He *could* try to persuade Patricia and Quentin to join him, but they won't. At least, it would take so much time and badgering and guilt-mongering on his part to persuade them, that it wouldn't be worth it.) It follows, then, that course of action *B* (which does not involve volunteering at Oxfam, but instead kidnapping, binding and gagging them) does not maximize utility. In terms of these assumptions, Denyer's argument can be presented formally as follows:

5. *D* can do *A* or *B*, but not both. [H]
6. *D* maximizes utility if and only if *D* does *A*. [H]
7. If *D* does *A*, then *P* and *Q* will potter. [H]
8. If *P* and *Q* potter, they will produce a total of 20 units of utility each. [H]
9. *P* and *Q* could do something other than pottering that would produce a total of 22 units of utility. [H]
10. If (8) and (9), then, if *P* and *Q* potter, they will fail to maximize utility. [Definition of "fail to maximize utility"]
11. If *P* and *Q* potter, they will fail to maximize utility. [8, 9, 10 Add., MP]

¹⁸See Denyer, 39–57.

¹⁹Nothing hangs on this simplification; if the argument works with twelve, it works just as well with two.

12. If *D* does *A*, then *P* and *Q* will fail to maximize utility. [7, 11, HS]
13. If *D* does *A*, then two failures to maximize utility will occur. [6, 12]
14. If *D* does *B*, then one failure to maximize utility will occur. [5, 6]
15. According to utilitarianism, *D* should do *A*. [6, Definition of Utilitarianism]
16. According to utilitarianism, *D* should do *A*, even though doing *A* does not minimize the overall number of failures to maximize utility. [13, 14, 15]
17. But if (16), then utilitarianism contains an ARC.
18. Therefore, utilitarianism contains at least one ARC. [16, 17 MP]

I believe this argument is unsound and reflects either a false conception of ARCs or a false conception of the ultimate objective of utilitarianism, or both. There are two ways of getting at this falsity: one through Denyer's description of his scenario, the other through the above analysis of his argument. Consider the scenario first: any clear-headed utilitarian will cry foul at Denyer's depiction of how utility works here, and how the utilitarian reasons about it. Denyer talks as if, on Saturday morning, when he thinks about the courses of action open to him, his actions and their utility are independent of everyone else's actions. But it is not so: according to utilitarianism, our job is to act so as to bring about the best future world, where this world includes other people *and* their actions *and* the effects of their actions, as well as our own, as these unfold into the future. It follows that, on Saturday morning, when he evaluates possible courses of action and concludes that one of them will maximize utility, he must take into account its effect on Patricia and Quentin and how it will combine with their likely actions to affect global utility in the future.

There are several ways, of course, by which Denyer might take his neighbors into account, some more direct and involved, others less so. For example, if he lived quite cut off from them, Denyer *could* take them into account by reasoning that, in the absence of any special reason to think otherwise, his actions and theirs would likely have little effect on each other, that they were causally independent. Or, given his epistemic limitations and the dangers of ignorant meddling, even if they are not causally independent, the least harmful option might be for him to treat them as if they *were* independent. Finally, if he is not so cut off from them, and he knows this, he may take them and their actions into account directly, and this he appears to have done. He does this when he reasons that Patricia and Quentin are unlikely to make their optimal contribution to global utility if left alone but any attempt by him to get them to do better is likely to backfire. If he could painlessly get them to join him in volunteering, he would have an obligation to do so, but Denyer's scenario rules that out, when it stipulates that for Denyer the only options are optimal volunteering or suboptimal kidnapping.

Whichever way he does it, when the Denyer in the story evaluates his options, he does so in the light of global utility, which includes their effects on Patricia and

Quentin and on their actions, and their actions' effect on global utility. He does this—or at least the utilitarian does this—because *his* concern is to maximize global utility, and nothing else. His concern is that *he himself* keeps global utility as high as possible, and not that his neighbors do so, except secondarily and contingently as this bears on *his* actions. A true utilitarian will see it as his job to choose actions and futures according to *whether* utility is maximized, not according to *how many* other people in those futures maximize it.

The significance of this point becomes clearer if we scrutinize my formal analysis of Denyer's argument. The problem with the argument so formulated is with premise (17), which (in expanded form) reads as follows:

17. If utilitarianism says that Denyer should do *A*, even though doing *A* does not minimize the overall number of failures to maximize utility, then utilitarianism contains an ARC.

Remember, Denyer wants to show that utilitarianism also is committed to ARCs, and premise (17) is where he claims that “to hold it wrong to fail to maximize utility, even though by committing one such wrong oneself, one could prevent a dozen such wrongs committed by others,” is to acknowledge an ARC. But that's false; it may on the surface *look* like an ARC, but it isn't one. Simply containing an “even though” construction does not suffice to make a judgment an ARC. We can see this if we consider the following judgments:

19. I must avoid great pain for myself in the future, even though this means accepting some pain for myself now.

20. I must pursue the greatest happiness for the greatest number, even though this requires self-sacrifice on my part.

21. I must help Patricia, even though this means harming Quentin.

22. I must do the best I can, even though this means doing repugnant things, sacrificing my projects, or violating my integrity as an agent.

23. I must seek to maximize utility overall, even though this sometimes means not consciously seeking to maximize utility at all.

All of these judgments contain an “even though” construction, but none of them is an ARC. Moreover, all of them are (or can be) perfectly acceptable to the clear-headed act utilitarian. In case this is not obvious, we need only view each judgment from the other side, so to speak, as an “in order to” construction. That is, the above are (respectively) equivalent to:

19.' I must accept some pain for myself now *in order to* avoid great pain for myself in the future.

20.' I must sacrifice my self *in order to* achieve the greatest happiness for the greatest number.

21. ' I must harm Quentin *in order to* help Patricia.

22. ' I must occasionally do repugnant things, sacrifice my projects or violate my integrity as an agent *in order to* do the best I can.

23. ' I must sometimes *not* consciously seek to maximize utility at all *in order to* maximize utility.

In each of these cases, some cost or harm is accepted (or imposed), but in each case, the cost or the harm is accepted for the sake of the ultimate utilitarian end, not despite it. And this ultimate utilitarian end—maximal utility—is supposed by the utilitarian to provide agent-neutral reasons for doing so, since, as Kagan observes, “. . . consequentialism can be usefully viewed as the theory that the *only* reasons for action are agent-neutral ones.”²⁰ Furthermore, as an end, *maximal utility* is distinct from the end of *some particular agent(s) acting so as to maximize utility*. While there is (according to the utilitarian) agent-neutral reason to desire the former, the utilitarian need acknowledge no agent neutral reason to desire the latter, except derivatively and contingently, when it is a necessary means to the former.

All of this bears on our original question as follows: absolutism is committed to ARCs, because it holds that for a certain kind of bad act, Φ , I mustn't Φ , even if by Φ -ing, I would minimize everyone's Φ -ing overall. That is, even if there is an agent-neutral reason to Φ , I have a contrary agent-relative reason *not* to Φ . Denyer wants his scenario to show that the utilitarian obligation to maximize utility entails exactly the same thing, but in that scenario, as in (19)–(23), I have no such agent-relative reason, because the only reasons involved are *ex hypothesis* agent-neutral ones. Thus the scenario does not show that utilitarianism (or any other sort of consequentialism) contains an ARC. The defense of absolutism as against consequentialism must move to other grounds.²¹

²⁰Kagan, 245.

²¹I thank Westmont College and the Centre de Recherche en Ethique de l'Université de Montréal (CREUM) for financial support during this writing of this paper, and Kelly Clark, Christopher Coope, Brad Hooker, David McNaughton, Ed Song, Jim Taylor and Stephen Zylstra for helpful written comments.