

## Consequentialism and Its Critics

Richard Arneson

word count 5785

{Final version published in Kelly Becker & Iain Thomson, eds., *The Cambridge History of Philosophy 1945-2015*, Cambridge University Press, 2019.}

In the second half of the twentieth century, the doctrine that came to be known as consequentialism has both taken hard thumpings and shown resilience in rebounding from attacks. In this process consequentialism has come to be fairly well understood, but elaborating and clarifying the varieties of nonconsequentialism are still works in progress, so it is perhaps premature to venture any definitive assessment on this dispute. Whether it is reasonable to accept a moral theory depends on the merits of its rivals. Perhaps any theory will have counterintuitive implications; perhaps the true one is least counterintuitive all things considered.

For purposes of this chapter *consequentialism* shall be understood as *act consequentialism*: One ought always to do an act that brings about an outcome no worse than what would have been brought about by anything else one might instead have done. In the comparison of the consequences of alternative acts available to the agent that determines which one she should do, omitting to act or doing nothing is counted as just another candidate act alongside others. Different standards of outcome assessment yield different versions of consequentialism. (Toward the very end of the chapter an alternative version of consequentialism, known as “rule consequentialism,” is briefly considered.)<sup>1</sup>

The formulation just indicated says that the permissibility status of an act is determined by its actual consequences. But at the time of decision, these may be unknown and perhaps unknowable. An alternative formulation of consequentialism says one should do what maximizes reasonably expected consequences.<sup>2</sup> In this chapter I set this issue aside. The criticisms to be considered apply will apply to either formulation of the doctrine.

The outcomes or consequences of acts have to be construed as overall long-term outcomes. If the outcome of a possible act would be the building of a beautiful skyscraper but one that ten years later will collapse and kill many people, the collapse is part of the outcome.

The justification for consequentialism is ready to hand and simple.<sup>3</sup> The idea is that rationality in conduct is maximizing—achieving the agent’s goals to the greatest possible extent. The question then becomes, what goals should the rational agent pursue. Different versions of consequentialism specify different goals. Consequentialism as such just imposes the constraint of impartiality on the rational agent’s goals: one person’s doing or getting something has exactly the same value, for purposes of deciding what action among those that the agent could do is morally required, as any other person’s doing or getting a relevantly identical thing.

What if there are several disparate impartial goals and they are not fully commensurable? Someone committed to consequentialism is not thereby committed to denying that possibility. Still, one ought to do what would bring about an outcome no worse, impartially assessed, than the outcome of anything else one might instead have done.<sup>4</sup>

Consequentialism identifies a family of views. The best-known member of the family is utilitarianism, which is act consequentialism plus this standard of outcome assessment: one outcome is better than another just in case the sum of individual welfare

contained in one is greater than the sum of welfare in the other. Individual welfare is summed across persons and, perhaps at a discount, other animals.

Versions of consequentialism can accommodate some prominent criticisms of utilitarianism. Utilitarianism gives fair distribution no weight in the assessment of outcomes: if two situations contain equal aggregate welfare sums, how welfare is distributed across persons makes no difference at all to the assessment of the situations; for the utilitarian it's all the same if saints fare well and scoundrels fare badly or the reverse, and it is all the same if the best off people have everything and the worse off nothing or if instead welfare is more evenly spread across persons.<sup>5</sup> Some might reject utilitarianism on the ground that the degree of fulfillment of individual moral rights either has some weight in the assessment of consequences or entirely determines whether one situation is morally better than another. If rights fulfillment is a goal to be promoted, it can figure in the consequentialist outcome assessment standard.<sup>6</sup>

Figuring out what is the most plausible version of consequentialism may not be enough to discover a plausible moral theory, let alone the most plausible one. The rival theories unite in denying that good consequences are all that matter in determining what is morally required, forbidden, or permissible to do.

In the following sections we rehearse several objections to consequentialism raised by prominent critics.

#### TOO DESTRUCTIVE OF INTEGRITY?

In 1973 Bernard Williams set off a fireworks of brilliant objections against consequentialism. These appeared in his "A Critique of Utilitarianism," but as he made clear, the target of his critique was consequentialism.

One central objection or series of objections revolved around the idea of integrity. Williams observes that an agent "is identified with his actions as flowing from projects and attitudes which in some cases he takes seriously at the deepest level, as what his life is about. . . . It is absurd to demand of such a man, when the sums come in from the utility network which the projects of others have in part determined, that he should just step aside from his own project and decision and acknowledge the decision which utilitarian calculation requires. It is to alienate him in a real sense from his actions and the source of his actions in his own convictions. . . . It is this, in the most literal sense, an attack on his integrity."<sup>7</sup>

Some problem is highlighted here, but exactly what is it? Suppose we stipulate that a person of integrity stands fast by her deepest convictions and values. A flat-footed response to Williams is that consequentialism does not attack the value of integrity, whatever that would mean, but instead embodies a proposal as to what are the values and convictions to which the person of integrity should be committed and to which, if he commits, he will stand fast.

This response brushes off the problem too quickly. As stated, consequentialism is a criterion of right action, not a decision making guide.<sup>8</sup> As such, consequentialism does not say what commitments agents should make. And part of what Williams is suggesting is that a commitment to bringing about the best is incompatible with bringing about much good, let alone the best.<sup>9</sup> Consider this argument:

1. If good consequences are maximized, people are committed to personal projects.

2. If people are committed to personal projects, they are disposed to pursue their personal projects even when doing so is not good-consequence-maximizing.

3. If people are disposed to pursue their personal projects even when doing so is not good-consequence-maximizing, they are not disposed to act in conformity with consequentialism.

4. If good consequences are maximized, people are not disposed to act in conformity with consequentialism.

The conclusion 4 follows logically from 1-3. Premise 1 is an empirical claim and premises 2 and 3 are partial explications of what it is to be committed to a personal project and to be disposed to act in conformity with act-consequentialism. Premise 1 could be true or false, but it surely could be true. Being committed to personal projects, in a way that rules out wholehearted commitment to act consequentialism, might be productive of better consequences than commitment to consequentialism. For example, loving particular types of things like skiing and playing jazz might bring about lots of skiing and jazz value into one's own life and the lives of people with whom one interacts, and also make one disposed to do wrong acts by the consequentialist standard. But the gains of the former might outweigh the losses of the latter, even if one actually does some wrong acts by that standard.

None of 1-4 is incompatible with consequentialism, the claim that one ought always to do whatever would bring about best consequences. But now some of what Williams is asserting comes into clearer view. The consequentialist morality makes the demand on the individual that she conform her conduct to this norm, and this demand alienates the individual from her deepest commitments. The alienation is a split or internal opposition in one's motivations: one cannot be wholeheartedly committed to one's personal projects and also recognize that one ought to do whatever would bring about the best, because the sincere recognition of what morality requires must induce some motivation to comply.<sup>10</sup>

Any morality might issue in demands on individual conduct that in particular circumstances conflict with what individuals would have to do to satisfy their deepest or strongest personal commitments. Sally's personal commitment might be to run a marathon fast, which it turns out she cannot do unless she murders Tom, but murdering Tom violates his moral rights and according to deontological morality is strictly prohibited. Also, the extent of conflict between moral demands and one's personal commitments depends on the substance of one's commitments (suppose one has a personal project of conquering all peoples). However, suppose that under scrutiny, it turns out that consequentialism, as it requires maximizing, in standard or normal circumstances makes demands that conflict to a far greater extent with the range of personal commitments that individuals in modern societies are likely to make, than the demands issued by rival moralities. So here's one objection against consequentialism latent in Williams's comments about integrity: Consequentialism is too demanding.

But alongside the "too demanding" objection is another distinct criticism: acceptance of consequentialism and conformity to it would be deeply alienating and as such destructive of value in our lives.

## TOO ALIENATING?

Peter Railton writes, “Living up to the demands of morality may bring with it alienation--from one’s personal commitments, from one’s feelings or sentiments, from other people, or even from morality itself.”<sup>11</sup> Living an unalienated life would then involve living in such a way that all of these elements are in complete harmony. One acts from personal projects one endorses and desires to fulfill, that are fully compatible with the moral requirements and aspirations one judges to be correct and is disposed to satisfy, having feelings and sentiments that resonate entirely positively with these projects and commitments, and relating to people in ways that are entirely free of conflict.<sup>12</sup> One’s actions are wholehearted, unaccompanied by doubt or hesitation.

Railton emphasizes another type of destructive conflict of motivation. The motives that induce one to take particular acts, even acts that are morally right and chime in with one’s personal commitments, may reduce the value gained by the acts. Example: One performs a service for a friend, or makes love to a romantic partner, from duty rather than the appropriate motive for the act—in the first case, affection for one’s friend, in the second case, sexual desire to make love to that particular person.<sup>13</sup>

If consequentialism required one always to act from the motive of seeking to bring about the best attainable outcome impartially assessed, adherence to consequentialism would be a giant wet blanket smothering and diminishing many sources of great value in life. But to reiterate a point already noted, consequentialism is a criterion of right action--not necessarily a decision making guide, and not necessarily the motive that causes one to act as one does.

Consider an example: if friendship is valuable, acts of forming and sustaining friendship can be right actions according to consequentialism. This can be so even if forming a friendship involves becoming disposed specially to favor one’s friendship partner, and this motivation may predictably lead one sometimes to perform acts that are wrong according to consequentialism.

The commitment to friendship cannot be wholehearted and unconstrained, if one is going to live a life that achieves as much by way of impartially good consequences as one feasibly can. Helping one’s friend in some circumstances might trigger large negative consequences, and one must be disposed to be alert to such situations and able to overcome friendship impulses to avoid pulling the trigger.

The person seeking to lead a life that scores high by consequentialist assessment looks to be inevitably up to her neck in alienation and internal conflict.<sup>14</sup> She is alienated from her personal projects and commitments, since pursuing them on some occasions will be seriously wrong according to consequentialism, and alienated from her friends and friendships, for the same reason, and alienated from whatever current plan for living her life she is now following, since it might turn out that following the plan will carry her over some cliff and down to very bad consequences.

But this may be OK all things considered. After all, a recipe for reducing all-sided alienation would be to minimize one’s commitments and feelings and beliefs and endorsements to a minimum, reducing as far as possible the possibility of conflict among these elements of one’s psychological makeup. This would be a poor strategy for attaining good for self and others. The consequentialist will say the ideal amount and character of alienation in one’s life is that which maximizes good consequences over the long run of one’s life.<sup>15</sup>

This resolution of the alienation concern may not satisfy the critics. Be that as it may, there's more to say in the Railton spirit. Railton defends consequentialism by insisting it is a criterion of right action and that's it. It need not and probably should not become the ever-present dominating motivation in an agent's life. It should instead play a background regulative role. It should be largely indirectly action-guiding. R. M. Hare appeals to a similar line of thought in wrestling with the objection that consequentialism is "too permissive" in allowing and requiring agents to violate deeply rooted basic moral constraints.<sup>16</sup> See discussion below.

#### TOO DEMANDING?

The "too demanding" objection might seem immediately to be a knockdown objection. Suppose that Sally is in a restaurant and can choose only between fish and chicken tacos, the choice affects no one other than herself, and the only relevant effect of her choice is that the fish tacos would be slightly tastier. Suppose other things equal, an outcome containing more enjoyment is better than one with less. So the best reachable outcome would be attained by choosing fish tacos and any other choice is morally wrong according to act consequentialism. This seems absurd.

One can mitigate this objection by accepting scalar consequentialism.<sup>17</sup> This says that any act that leads to a less than best outcome is wrong, but acts are more wrong, the greater the shortfall between the moral value of the best outcome one could have reached and that of the outcome of what one actually did. If the shortfall is tiny, one's act is wrong, but barely wrong, and it would be wrong to fret about the shortfall or blame oneself for doing what is after all barely wrong. In the example Sally's choice of chicken tacos is just barely wrong. Arguably that verdict delivered by a plausible consequentialism is not counterintuitive.

This response even if deemed acceptable does not significantly dent the "too demanding" objection, because in actual and likely future circumstances, with the Earth beset by widespread poverty, misery, war, and premature death, and mechanisms in place that enable many of us to contribute to alleviating these ills with reasonable effectiveness, and keep giving at cumulative enormous cost to ourselves, the shortfall between the actual consequences of what one does and the best outcome that might instead have been reached is huge.

Samuel Scheffler addresses this issue in 1982. He is explicitly responding to Williams's arguments, and working also to accommodate ideas advanced by Thomas Nagel.<sup>18</sup> Each person sees things from a unique personal perspective, and from this perspective there are things that matter to that person, but which do not show up as reasons for action at an impartial level, where a reason for any person is a reason for all persons. An adequate morality cannot simply endorse impartial reasons alone but must somehow split the difference between the personal and impersonal perspectives. How to do this is not clear.

Scheffler suggests that morality gives each agent a personal prerogative to weight her own personal, agent-relative concerns more heavily than other people's in deciding what it would be best to do. She is then morally bound when acting only to produce as much impartially assessed good as would have been achieved had she acted on her personally weighted calculation. Each person in acting then always has a personal

prerogative to do what brings about less than the best up to some shortfall. This is a limit that varies depending on what is at stake.

A stronger personal prerogative what would place a flat upper limit on the amount of sacrifice of what matters to her that any individual is morally obligated to incur, in each choice she makes or over stretches of her life or over her whole life. One could couple this with a requirement that the sacrifice one is bound go incur, one must expend efficiently toward advancing the greater impartial good.

As stated, the Scheffler prerogative permits one to do less than best to achieve any aim or none, so long as the acceptable shortfall earned by one's interests in the situation at hand is not exceeded. One is permitted not to save the city so that one is morally at liberty to pursue an important project, but one is further at liberty to pursue any trivial aim of little or no personal concern so long as the prerogative limit is not exceeded.

Another type of prerogative would limit the sacrifice one is required in order to bring about overall greater good in some situation by what would be a fair share if all who are duty bound to help contributed a similar fair share. The idea is that would be an objectionable implication of a proposed principle of beneficence if it required agents to do more and make greater sacrifices when other people are failing to conform to beneficence requirements. "Why should I have to take up the slack when others fail to do their fair share?," one might complain.<sup>19</sup> A problem with this suggestion is that the unfairness among potential contributors if someone was required to take up the slack and do more than her fair share would seem generally to be outweighed by the unfairness of the persons in peril or misery to whom help is owed not getting the full amount they should receive.

The Scheffler personal prerogative idea rests on the assumption there are agent-relative reasons, genuine reasons that give the agent a reason for action but that do not generate reasons for others. The consequentialist will oppose the assumption and contend that if some one has a reason to do something, others thereby are given reason to promote or facilitate the doing.

The Scheffler personal prerogative at least in the form in which Scheffler endorsed it is a limit on an assumed background duty of beneficence, a duty to improve the world, bring about good outcomes according to an impartial standard of what makes outcomes better or worse. What Scheffler shares with the consequentialism he opposes is considerable. The line between those who accept some such duty of beneficence as an important component of morality and those (such as Nozick) who deny any such duty and regard beneficence as always morally optional is perhaps more important in the end than the dispute between those who assert, and those who deny, that the duty of beneficence must be balanced against, and limited by, a moral freedom that persons have to live as they choose, to some extent.

### TOO PERMISSIVE?

Besides being branded as too demanding, consequentialism is also faulted for being too permissive. The thought is that a person is not morally at liberty to do whatever would bring about greater good when doing so would either run roughshod over bystanders in ways that harm them or use the opportunity provided by the presence of others to advance one's ends in ways that harm them. There are moral constraints that prohibit us from always doing whatever would bring about the best outcomes.

At the start of his great work *A Theory of Justice*, John Rawls articulates a morality of constraint: “Each person possesses an inviolability founded on justice that even the welfare of society as a whole cannot override. For this reason justice denies that the loss of freedom for some is made right by a greater good shared by others. It does not allow that the sacrifices imposed on a few are outweighed by the larger sum of advantages enjoyed by many.”<sup>20</sup> *Justice* here refers to enforceable moral duties, those that justify coercion if needed to secure compliance. Rawls surmises that utilitarianism can be understood as based on the idea that the principle of rational choice for a single individual can be extended across society. Just as it is rational for an individual to impose on himself the lesser pain of a trip to the dentist to avoid a later greater pain of serious toothache, it is rational to impose on one individual a loss equivalent to the dentist trip pain in order to gain for another the greater benefit equivalent to serious toothache avoidance. Disagreeing with this idea, Rawls observes, “Utilitarianism does not take seriously the distinction between persons.”<sup>21</sup> He could have said the same about consequentialism.

Accepting the idea that morality at the level of fundamental principle includes constraints, one could also hold that imposing sacrifice on some for the sake of aiding others could be acceptable in those cases in which the beneficiaries have a moral right to be benefited and the others, on whom sacrifice is imposed, are under a moral duty to aid them. Also, no matter how morally bad murdering a person is, two relevantly similar acts if murdering would be worse, so a morality of constraints as so far described might yet allow the bringing about of best outcomes by doing whatever will bring about the greatest overall fulfillment (or equivalently, least overall violations) of moral rights.

Robert Nozick closes off these possibilities in a discussion that sharpens and clarifies the issues. He suggests that our ordinary common-sense idea of a moral right should be understood as a side constraint on choice rather than a goal to be fulfilled. He also makes a substantive proposal about the content of basic moral rights: they are one and all negative rights not to be harmed or subjected to interference not positive rights to be benefited.<sup>22</sup>

The side constraint idea says that moral rights should enter the determination of what is permissible as follows: from the set of possible acts one could now perform, strike out all those that would violate anyone’s rights and limit your choice of action to those that remain. This implies one should not steal from Pete in violation of one of his moral rights even if doing so would prevent worse rights violations by other people.

If people had rights to be helped as well as rights not to be harmed, there could be cases of conflict, in which one could fulfill one’s duty to help only by acting against one’s duty not to harm, and in these cases perhaps one should weigh what is at stake, and sometimes rightly judge that the duty to help outweighs the duty not to harm. Rejecting all positive rights and their corresponding duties, Nozick asserts that we should countenance no such tradeoff and moral balancing scenarios.

But suppose we should reject as fanatical the claim that rights are exceptionless so that one should never do what a moral right forbids. That concession still leaves intact a robust doctrine that opposes consequentialism by affirming the existence of moral constraints that should be obeyed even if acting against them would bring about a greater sum of good, at least if the sum is below some threshold, the location of the threshold depending on the seriousness of the constraint.

In the face of these objections, the consequentialist has a strategy of response. How successfully it meets the objections is left for the reader to decide.

The strategy invokes an error theory. The consequentialist can explain how people might find consequentialism counterintuitive and problematic in its recommendations for conduct, even if consequentialism is true. They would continue to find consequentialism dubious even in a society that was well regulated to educate and socialize people in ways that are as conducive to conformity to consequentialism as is feasible.

The explanation is suggested or at least hinted at by many earlier thinkers but received full articulation in the discussion of “The Archangel and the Prole” in R.M. Hare’s 1981 book *Moral Thinking*.<sup>23</sup> Hare notes that human persons are not angels, but rather tend to be selfish, not well informed about facts that are relevant to choices they face, and only imperfectly intelligent, so not very good at integrating such factual knowledge as they do possess into sensible calculation of what actions and policies to choose. One could add that they are imperfect reasoners about normative issues. Moreover, different persons vary in the degree to which they have these various rational decision making deficits, and any given person will be afflicted variously along each of the dimensions of deficit in different contexts and settings.

Given these background conditions, the decision procedure of considering which of the alternative acts one could perform would bring about best consequences, on each occasion of choice, would predictably be disastrous, unless one happens to have archangelic powers of reasoning on this occasion. For most of us most of the time, a better decision procedure would be rather rigidly to follow common-sense moral prohibitions and requirements such as Keep your promises, Don’t steal, Tell the truth, and Don’t harm innocent nonthreatening persons. These common-sense rules will incorporate familiar constraints against harming others and using them to their disadvantage to advance one’s purposes. Moreover, for the rules to function well, we must be trained to internalize them, that is, to treat their dictates as intuitions of conscience, which we have reason to follow.

The resulting position is odd. If we are well trained, we will believe, for example, that we ought to refrain from aggressively physically attacking others. We will believe that these and other norms give us basic non-derivative reasons for choice and action. But if we are reflective, we will also believe that only the fundamental level consequentialist principle gives basic non-derivative reasons for action, and the norms at the level of common-sense moral rules (and similarly for laws and other social norms) are just tools for helping to bring about greater conformity by ourselves and others to consequentialism. Roughly, we believe not that killing the innocent and lying and so on are intrinsically wrong but that better consequences will be brought about by pretending that this is so. Our moral beliefs are in tension, to put it mildly.

The common-sense rules we should follow are the ones actually in place in society and actually coordinating people’s expectations of what others are doing. But our allegiance to the existing code should vary in strength depending on how far it deviates from the code that in actual circumstances would be both feasible and best promote good consequences. Moreover, we should engage in acts of pushing for code improvement whenever such acts would be consequentially justified. Furthermore, depending on the degree to which humans vary in their resemblance to archangels (perfect moral decision



makers) or proles (perfectly bad moral decision makers), we should train and encourage people including ourselves to follow archangelic fully deliberative decision procedures or instead to follow the prolish decision procedure of mechanically following prevailing moral rules.

None of this amounts to a reason to believe that consequentialism is correct. The position rather is that if there is a convincing case for the truth of consequentialism, we should take at a discount our nonconsequentialist moral intuitions arising from our internalization of common-sense moral rules prevailing in our society. After all, we would have these or similar intuitive judgments even if consequentialism was best supported by the reasons there are.

The idea that moral thinking operates at different levels such as the theoretical and the intuitive level, no doubt is disturbing and counterintuitive. However, not only consequentialism, but any nonconsequentialist morality that includes moral concern for the degree to which its dictates are followed will accept the need for different levels of moral thinking. Given that we are not fully archangelic but instead to varying degrees prolish, the correct nonconsequentialist principles may be too complex or otherwise unsuitable to serve as a practical decision making guide. Hence we should distinguish the nonconsequentialist criterion of right action and the best decision procedures we can devise and follow to bring about our conformity with fundamental nonconsequentialist morality. A sensible nonconsequentialist morality will then need to distinguish different levels of moral theory, the derivative levels being tools to facilitate greater conformity to fundamental level principles. The problems we encounter in working out coherent and sensible multi-level theory will be shared by consequentialists and nonconsequentialists so cannot be claimed to be any reason for rejecting consequentialism and opting instead for nonconsequentialism.

#### A FALSE SHORT CUT, A RED HERRING, AND A LONG MARCH

In 1992 Frances Kamm suggested a way to break the deadlock of intuitions between consequentialists and nonconsequentialists on the moral reason-giving force of moral constraints by appeal to the status of the human person as a rational agent and therefore inviolable to a high degree.<sup>24</sup> Being inviolable to a high degree, a person is not permissibly subject to being harmed either by actions that cause greater good or as a byproduct of such actions. This high status is a value that inheres in each person and calls for respect. Thomas Nagel and (some years later) Michael Otsuka have endorsed this appeal. In Otsuka's words, "nonconsequentialist prohibitions are justified because they properly reflect our elevated moral status."<sup>25</sup>

But this move is offset by another.<sup>26</sup> The consequentialist can appeal to the status of the human person as a rational agent and therefore unignorable to a high degree. Being unignorable to a high degree, each person's interests give rise to reasons for their fulfillment by everybody and anybody that are not diminished in strength by the sheer fact that their fulfillment would have to run through a causal path involving losses to others' interests.

The nonconsequentialist appeals to the thought that by possessing rational agency capacity some entities qualify as persons and have a special dignity and worth.<sup>27</sup> The nonconsequentialist associates this dignity with the status of inviolability, the consequentialist associates this dignity with the status of unignorability. Which interpretation is better? We have come back to the issue, whether morality forbids or

requires harming whenever harming would bring about greater good. The appeal-to-status gambit has not made progress toward resolving the issue.

Consequentialism tends to be associated with aggregation: the ideas (1) that for any harm, however bad, that falls on a person, there is some perhaps very large number of tiny harms to many persons that is worse, and (2) that for any moral wrong, however horrible, that anyone does, there is some perhaps very large number of tiny wrongs that is overall morally worse. For example, could some very large number of barely noticeable mild headaches suffered by a very large number of people add up to amount of disvalue greater than the horrific badness of the cutting short by death of the life of a single child who otherwise would have lived a long life at a high level of flourishing? Could a very large number of trivial moral wrongs such as thefts of a penny in aggregate be overall morally worse than a single horrific torture-murder? Those who regard consequentialism as committed to aggregation tend to cite this as a defect in consequentialism. But this is a red herring. There are versions of consequentialism that reject aggregation and accept that some goods and harms have lexical priority over others. That is, a consequentialist view could accept an outcome assessment standard according to which the slightest amount of some excellent good outweighs any quantity of lesser-quality goods—so no number of licks of ice cream could yield as much value as a sublime moment of appreciating the best poetry. There are versions of nonconsequentialism that accept aggregation. For example, one might hold that one should respect individual moral rights even when violating rights would yield a somewhat better outcome, but also hold that for any right, there is (possibly enormous) some amount of harm that would accrue to nonrightsholders if the right is respected, by comparison with the (very small) amount of harm that would accrue to rightsholders if the right is violated, such that in these circumstances the right gives way and one ought to do what the right forbids.

Moreover, when we consider the issue, aggregation, the idea that many small goods can outweigh a large good, and that rights and wrongs trade off against each other in the determination of what ought to be done, is hard to resist without introducing rank implausibility somewhere in one's morality.<sup>28</sup>

What would be useful is the further development of work now ongoing toward characterizing the nature and content of plausible individual moral rights with their corresponding duties owed to rightsholders. This involves clarifying the distinctions partially constitutive of these rights such as between doing and allowing and between intended and foreseen but unintended consequences of what we do and allow and between harmful agency that does and does not benefit from the presence of those harmed.<sup>29</sup>

#### ACT CONSEQUENTIALISM OR RULE CONSEQUENTIALISM?

Throughout the second half of the twentieth century, there have been recurrent efforts to resolve the nagging intuitive conflict between what ordinary common-sense moral rules permit and require and the revisionary implications of the principle that one ought always to do whatever would bring about the best reachable outcome.<sup>30</sup> Some consequentialists have played down the conflict by speculating that in actual circumstances these two independent sets of norms will tend to converge in their practical implications. But it's doubtful this speculation is true, and anyway, the moral doctrine worthy of our acceptance should yield acceptable judgments for what to do in possible

circumstances we can imagine. I have already described a consequentialist strategy that consists in explaining away the appearance that common sense moral norms are genuinely normative. Another complementary strategy argues directly on moral grounds that the “common-sense” rules are revealed under scrutiny to be a mess and to lack authority.<sup>31</sup>

Another strategy with perennial appeal has been to hold that the moral rules we should accept are the ones whose general acceptance would have best consequences and that at the level of individual choice of conduct one ought to follow those rules. If the moral rules selected by this principle are pretty close to the already accepted common-sense moral rules that seem to have some claim on our allegiance, the consequentialism-versus-ordinary-rules conflict diminishes or disappears. Or so it might be thought.

Many questions arise. If the ideal code is not actually established and accepted, why should the fact that if it were accepted, that would have good consequences, render it the case that here and now we should conform our behavior to the code? If the code yields implications for conduct opposed to consequentialism, why do what in hypothetical nonfactual circumstances would work for the best rather than what in present actual circumstances would work for the best?

Following and refining suggestions by Richard Brandt, Brad Hooker in 2000 proposed this version of rule consequentialism: “an act is morally permissible if, only if, and because the act is allowed by the code of rules whose internalization by the overwhelming majority of everyone in future generations would maximize aggregate expected welfare,” adjusted to give some priority for the worst off.<sup>32</sup> I won’t try to assess this interesting proposal except to note that the questions posed in the previous paragraph perhaps still have bite.

## CONCLUSION

In the past 75 or so years, progress has been made in the controversy between advocates of consequentialism and their critics. The progress has consisted in clarifying options, not in bringing about consensus among disputants.

---

<sup>1</sup> . There are forms of consequentialism different from the act and rule varieties. See, e.g., Adams 1976.

<sup>2</sup> . Pettit 1997 makes a case for identifying the consequentialist criterion of right action in terms of what it would be morally rational for an agent to do given available evidence. Parfit 2011 suggests there are just two different noncompeting notions here of what it is right to do.

<sup>3</sup> . Mill states it clearly (Mill 1863, chapter 1).

<sup>4</sup> . For fascinating discussion of the complexities that limited commensurability insinuates into the determination of what one ought to do (and of other complexities related to the assessment of consequentialism as well), see Hare 2013.

<sup>5</sup> . The idea that how the good is distributed across persons matters to the assessment of outcomes is asserted in Scheffler 1982. He affirms the more specific prioritarian idea that

the moral value of getting a benefit of a given size to a person is greater, the lower the person's well-being level absent this benefit. The idea that the moral value of a benefit's accruing to a saint is greater than its accruing to a scoundrel receives exhaustive treatment in Kagan 2012.

<sup>6</sup> . See Sen 1982.

<sup>7</sup> . Williams 1973, 116-117.

<sup>8</sup> . Bales 1971.

<sup>9</sup> . Versions of this argument appear in Williams 1972 and 1973.

<sup>10</sup> . Parfit 2011 takes this point to tell against the idea that rational persons would choose that everyone conform to (act) consequentialism.

<sup>11</sup> . Railton 1984,

<sup>12</sup> . Hurley 2009 urges that if the consequentialist could convince us that what one morally ought to do is always the act that would bring about best consequences impartially assessed, we should then acknowledge a deep conflict between morality and practical reason. What one would have most reason to do, all things considered, would often be to pursue one's deepest projects or follow one's own heart's desire.

<sup>13</sup> . Compare Adams 1976.

<sup>14</sup> . In this connection see Stocker 1976.

<sup>15</sup> . Railton 1984.

<sup>16</sup> . Hare 1981.

<sup>17</sup> . See Norcross 2006. See also Brink 2013 on Mill's anticipation of the doctrine.

<sup>18</sup> . Scheffler 1982, Nagel 1986.

<sup>19</sup> . Murphy 1993.

<sup>20</sup> . Rawls 1999, 3.

<sup>21</sup> . Rawls 1999, 24. For discussion, see Brink 1993.

<sup>22</sup> . Nozick 1974.

<sup>23</sup> . Hare 1981. For discussion of Henry Sidgwick's consideration of the degree to which the consequentialist criterion of right, in actual and likely circumstances, requires the teaching of a moral code much like conventional common-sense morality as a practical decision-making guide, see Hurka 2014, chapter 7. A classic defense of common-sense moral rules against utilitarianism (and consequentialism) is Ross 1930, also discussed in Hurka 2014. On Hare's strategy, see also Railton 1984 and Eggleston 2014.

<sup>24</sup> . Kamm 1992.

<sup>25</sup> . See Nagel 1995; Otsuka 2009, 66.

<sup>26</sup> . Lippert-Rasmussen 1996.

<sup>27</sup> . A clear articulation of this idea is in Nozick 1974, 48-49.

<sup>28</sup> . On aggregation regarding right action, see Thomson 1990 and Scanlon 1998. For doubts about the rejection of aggregation in the domain of right action see Parfit 2011 (chapter 21) and in the domain of value, Norcross 1997.

<sup>29</sup> . The literature on these topics is voluminous and impressive. See especially Nozick 1974, Thomson 1990, Kamm 1993 and 2007, Scanlon 1982 and 1998, Quinn 1989a and 1989b, Parfit 2011.

<sup>30</sup> . Lyons 1965 is a pathbreaking clarification of discussions of rule-utilitarianism. See also Rawls 1955 and for skepticism about rule-utilitarianism, Smart 1955 and 1973.

<sup>31</sup> . A notable exemplar of this strategy is Kagan 1989.

---

<sup>32</sup> . Brandt 1992, Hooker 2000 and 2014. The quote in the text is from Hooker 2000, 32. For criticism, see Arneson 2005 and Portmore 2009.