

UTILITARIANISM AND CONSEQUENTIALISM: THE BASICS

I. Hedonistic Act Utilitarianism (HAU)

A. Definitions

Hedonistic Act Utilitarianism: An act is morally permissible if and only if it maximizes hedonic utility, and if an act is morally permissible, then what makes it so is ultimately the fact that it maximizes hedonic utility.

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

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

B. Its Appeal

The thought is that HAU is just a natural extension of the familiar balancing and maximizing conception of rationality from diachronic intrapersonal contexts to diachronic interpersonal (moral) contexts.

The diachronic intrapersonal context: The temporal location of a benefit or harm within a life does not, as such, have any rational significance. We should, therefore, be temporally neutral with respect to *when* harms and benefits occur in our lives. And we should be willing to make sacrifices now for the sake of future benefits that will more than offset these sacrifices. Thus we should be willing to make trade offs (to balance) for the sake of maximizing our welfare over time.

The diachronic interpersonal (moral) context: From the moral point of view, no one person's welfare is any more or less important than another's. So, the location of a benefit or harm within a population of persons has no moral significance. And thus, as we move from the point of view of individual rationality to the point of view of morality, we should not only be temporally neutral with respect to *when* harms and benefits occur, but also be person-neutral (or impartial) with respect to *who* receives some harm or benefit. Furthermore, we should again be willing to make trade offs (to balance) for the sake of maximizing welfare over time and across individuals.

C. A Practice Quiz: Getting a Feel for the Theory

Answer the following questions according to what hedonistic act utilitarianism (HAU) entails—answer either “Necessarily True” or “Not Necessarily True.”

1. An act is morally impermissible if it causes someone a lot of pain and no pleasure.

Answer: Not necessarily true. An act that causes one person pain and many others pleasure may, in the aggregate, produce more hedonic utility than any other available act alternative would. Suppose, for instance, that there are exactly three act alternatives available to the agent: a_1 - a_3 . Further suppose that these acts will have effects on only three persons: p_1 - p_3 . Now let ‘ $HU(p_1)$ ’ stand for ‘ p_1 ’s hedonic utility’, and let ‘ HU ’ stand for ‘the total aggregate utility’. Suppose, then, that the scenario is this:

<i>act</i>	<i>HU(p₁)</i>	<i>HU(p₂)</i>	<i>HU(p₃)</i>	<i>HU</i>	<i>moral status</i>
a_1	+10	-100	+130	+40	permissible
a_2	+5	+5	+5	+15	impermissible
a_3	+15	+25	0	+40	permissible

Lesson: On HAU, it’s the aggregate that matters. This means that one may permissibly impose great sacrifices on some for the sake of greater gains in the aggregate.

2. An act is morally impermissible if it produces, in the aggregate, more pain than pleasure. In other words, an act is morally impermissible if it has negative hedonic utility.

Answer: Not necessarily true. If all the available act alternatives have negative hedonic utilities as well, then the given act will be permissible provided that there is no alternative with less negative hedonic utility. Consider, for instance, the following scenario:

<i>act</i>	<i>HU(p₁)</i>	<i>HU(p₂)</i>	<i>HU(p₃)</i>	<i>HU</i>	<i>moral status</i>
a_1	+10	-100	+30	-60	impermissible
a_2	+5	-5	-5	-5	impermissible
a_3	+15	+25	-41	-1	permissible

Lesson: According to HAU, whether or not an act is permissible can only be determined by comparing its hedonic utility with that of its alternatives.

3. If an act produces pleasure for many and pain for no one, then it is morally permissible.

Answer: Not necessarily true. On HAU, it's not always sufficient to produce pleasure for many and pain for no one; one must maximize hedonic utility, producing as much hedonic utility as possible. Consider, for instance, the following scenario:

<i>act</i>	<i>HU(p₁)</i>	<i>HU(p₂)</i>	<i>HU(p₃)</i>	<i>HU</i>	<i>moral status</i>
a ₁	+10	-100	+130	+40	impermissible
a ₂	+100	+100	+100	+300	impermissible
a ₃	+100	+101	+100	+301	permissible

Lesson: Because HAU requires that we always maximize hedonic utility, it will be a very demanding moral theory. It will require each of us to perform an altruistic self-sacrificing act whenever the self-sacrifice that we thereby make will be more than offset by altruistic benefits we thereby bestow on others.

4. It is morally impermissible to listen to *Metallica* on the car's CD player if the other two people in the car would get more pleasure from listening to *The Beatles*. Assume that listening to *Metallica* and listening to *The Beatles* are the only two available options. Assume that you and the other two people in the car are the only ones who will be affected by this choice.

Answer: Not necessarily true. If HAU was a democracy, it would be "one hedon/dolor, one vote"; it would not be "one person, one vote." HAU is not a form of majoritarianism. To illustrate, let a₁ = listen to *Metallica*, and let a₂ = listen to *The Beatles*. Now suppose that the scenario is this:

<i>act</i>	<i>HU(p₁)</i>	<i>HU(p₂)</i>	<i>HU(p₃)</i>	<i>HU</i>	<i>moral status</i>
a ₁	+10	+10	+80	+100	permissible
a ₂	+15	+15	+15	+45	impermissible

Lesson: Some people will loom larger than others in the moral calculus simply because they stand to gain or lose more than others do. There might even be *utility monsters*: people who get so much pleasure from getting what they want that the rest of us would all be morally required to do whatever we can to fulfill their every desire.

5. In an effort to help his patient, Dr. Smith injects Jones with a shot of penicillin after dutifully asking him whether he's allergic to penicillin and

hearing Jones say that he isn't. Nevertheless, it turns out, unbeknownst to Jones even, that Jones is severely allergic to penicillin, and, consequently, he dies as a result of the injection. Assuming that there would have been more hedonic utility had Dr. Smith not injected Jones with the shot of penicillin, it follows that what Dr. Smith did was wrong.

Answer: Necessarily true. On HAU, the agent's motives and intentions are irrelevant in assessing whether his or her act is permissible or not; it's only the consequences of the action and of their alternatives that matter.

Lesson: HAU says nothing about what sort of motive or intention we should have when we perform an act. Thus it does not require that we are always motivated to act out of a concern to maximize hedonic utility. Sometimes we will act out of a partial concern for ourselves or for our loved ones, but what's important is that our acts accord with HAU, not that they are motivated by a desire to act in accordance with HAU.

6. Anyone who acts wrongly as Dr. Smith did above is a morally bad person, deserving of both blame and punishment.

Answer: Not necessarily true. As it stands, HAU evaluates actions, not agents. So it entails nothing about whether such a person is morally bad or deserving of blame. HAU does entail that we should act so as to impose some punishment on someone or act so as to ascribe blame to that someone if and only if doing so would maximize hedonic utility, but it's not necessarily the case that blaming or punishing such a person would maximize hedonic utility.

Caveat: Even if HAU commits us to the view that an agent's motives and intentions are irrelevant when assessing the moral value of her actions, it doesn't commit us to the view that an agent's motives and intentions are irrelevant when assessing the moral goodness of her character.

7. It is morally permissible to torture an innocent baby in certain circumstances.

Answer: Necessarily true. On HAU, the ends justify the means, and thus virtually any type of act can, in principle, be justified and, indeed, will be justified if it maximizes hedonic utility. Let a_1 = torture the innocent baby, a_2 = refrain from torturing the innocent baby, p_1 = the would-be torturer, and p_2 = the would-be torture victim. Now consider the following scenario:

<i>act</i>	<i>HU(p₁)</i>	<i>HU(p₂)</i>	<i>HU</i>	<i>moral status</i>
------------	--------------------------	--------------------------	-----------	---------------------

a ₁	+100	-99	+1	permissible
a ₂	0	0	0	impermissible

Lessons: First, on HAU, everyone's welfare counts, and counts equally, regardless of desert. Second, on HAU, it is permissible to maim, murder, and even massacre innocent people provided that doing so will maximize hedonic utility. On HAU, most any type of act is, in principle, justifiable. Also, HAU leaves no room for the idea that people have rights that make it wrong to use them as a means to the greater good.

8. The effects that our actions will have on people's welfare millions of years from now are just as important as their immediate effects.

Answer: Necessarily true. On HAU, all effects matter no matter how remote in space or time. Let a₁ = stay after class and talk to John, a₂ = go home immediately after class, stopping to talk to no one, p_c = the six billion or so people that currently exist, and p_f = the ten billion or so people that will exist three generations from now. And let's suppose that if I stay after class and talk to John, John will after we're finished go to the local coffee shop and meet a woman standing behind him in line, a woman whom he falls in love with and has a child with. Assume that this child's grandchild ends up being an evil fascist dictator who causes untold suffering. Lastly, assume that had I not stayed after class to talk with John, he would have left for the coffee shop much earlier, and, consequently, this evil fascist dictator would have never existed. Suppose, then, that the scenario is this:

<i>act</i>	<i>HU(p_c)</i>	<i>HU(p_f)</i>	<i>HU</i>	<i>moral status</i>
a ₁	+1	-100 tr.	-999,999,999,999	impermissible
a ₂	0	0	0	permissible

Lessons: First, given that there are, at any given moment, an infinite number of available act alternatives, and given that one would need to know what effects each of these will have from now until the end of time to know what their hedonic utilities are, one can never know for certain what one morally ought to do on HAU. Some have, therefore, objected to HAU on the grounds that it is too impractical to be action guiding. Second, given that there is, at most every moment in everyone's life, some great tragedy that one could potentially advert but which one fails to advert due to a lack the relevant knowledge about the future consequences of one's available act alternatives, it follows that most everyone is, at most every moment, acting immorally. There is, at most every moment, some future tragedy that we could advert if only we did one thing rather than another.

9. While waiting at a crosswalk, Mr. Smith sees a stranger about step off the curb into the path on an oncoming, speeding bus. Without thinking, Mr. Smith reaches out, grabs the stranger, and pulls him back onto the curb, saving his life. Mr. Smith does this without even thinking about what the hedonic utility of this act or any of its alternatives might be. What Mr. Smith did was wrong.

Answer: Not necessarily true. On HAU, what matters is not how you came to act in the way that you did, but whether the way you in fact acted maximized hedonic utility. Let a_1 = your instinctually grabbing the stranger and pulling him back to safety, a_2 = your calculating the likely hedonic utility of each of the various act alternatives available to you while the stranger is struck and killed by the bus, p_1 = you, and p_2 = the stranger. Suppose, then, that the scenario is this:

<i>act</i>	<i>HU(p₁)</i>	<i>HU(p₂)</i>	<i>HU</i>	<i>moral status</i>
a_1	+20	0	+20	permissible
a_2	-20	-999,980	-1 million	impermissible

Lesson: HAU is a criterion of rightness (i.e., a description of what the fundamental right-making and wrong-making features of actions are), not a decision procedure (i.e., not a procedure for deciding what to do). To understand the difference, consider the difference between the criterion for being HIV infected (or brain dead), and the procedure that we might use to determine whether someone is HIV infected (or brain dead). And note that HAU is no more impractical than the view that the most prudent act is the one that will maximize one's utility over time. And no one would suggest that we ought to reject the investment principle "Buy low, and sell high" just because it is often times practically impossible to know whether the price of a stock is headed up or down.

10. If you promised a kid ten bucks to wash your car and he did as promised (he washed your car and did a nice job), you should give him the ten bucks.

Answer: Not necessarily true. Whether or not you should give the kid the ten bucks depends, not on whether you made a past promise to do so, but rather on whether your doing so would maximize hedonic utility.

Lesson: HAU is entirely forward-looking. What happened in the past is, on HAU, irrelevant in determining what one morally ought to do. The past is irrelevant, because nothing you do now can affect the hedonic utility of any past event.

11. Mr. Smith has an extra \$500 this month after paying all his bills. He uses that \$500 to buy his son a set of encyclopedias to help with his son's schoolwork. But had Smith instead used that \$500 to buy some stranger's kid a set of encyclopedias that would have produced slightly more hedonic utility. What Mr. Smith did was wrong.

Answer: Necessarily true. On HAU, everyone's pleasures and pains count equally, regardless of one's relationship to those affected. The hedonic utility of one's friends and family members are no more (or less) important, on HAU, than that of complete strangers.

Lesson: On HAU, there are no special obligations, and special relationships have no intrinsic importance.

12. In deciding what to eat for breakfast this morning, Mr. Johnson is morally required to make the choice that produces the greatest hedonic utility. Assume that, according to HAU, Mr. Johnson ought to eat something for breakfast this morning, and that what choice he makes will affect only himself.

Answer: Necessarily true.

Lesson: According to HAU, morality is pervasive. Every practical decision and choice is a moral one, even those that are quite trivial and that have no ramifications for anyone but the agent. There is no moral asymmetry between the self and others. You have just as much moral reason to promote your own hedonic utility as to promote anyone else's.

13. The King of Siam must adopt either policy X or policy Y. This decision will affect only those people belonging to one of the following two populations: population A and population B. Population A consists in a hundred poor persons with 10 hedons each. Population B consists in a hundred affluent persons with 100 hedons each. The King of Siam adopts policy X, which has no effect on the people in population A, but raises the number of hedons for each person in population B by 90 hedons. Had the King of Siam adopted policy Y instead, this would have had no effect on the people in population B, but would have raised the number of hedons for each person in population A by 90 hedons. Policy X, therefore, worsens the already existing inequality that exists between the two populations, whereas policy Y would make things equal by raising the level of welfare of the worse off. Given that the people in population A are equally deserving of 100 hedons as the people in population B are, it was wrong for the King of Siam to have adopted policy X.

Answer: Not necessarily true. Let a_1 = adopt policy X, a_2 = adopt policy Y, p_A = population A, and p_B = population B. The scenario is as follows:

<i>act</i>	<i>HU(p_A)</i>	<i>HU(p_B)</i>	<i>HU</i>	<i>moral status</i>
a_1	+10,000	+10,000	+20,000	permissible
a_2	+1000	+19,000	+20,000	permissible

Lesson: On HAU, equality is of no intrinsic importance. HAU is not sensitive to concerns about desert or distributive justice.

14. On HAU, there are no supererogatory acts (i.e., acts that go above and beyond what duty requires).

Answer: The answer depends on what the correct definition of 'supererogatory act' is. Some define supererogatory acts as those that do more for the sake of others than duty requires. Others define supererogatory acts such that S's performing x is supererogatory only if there exists some available alternative, y , such that (a) S is morally permitted both to perform x and to perform y , and (b) S has more moral reason to perform x than to perform y . If the former definition is correct, then HAU can accommodate supererogatory acts. But if the latter is the correct definition, then HAU cannot accommodate supererogatory acts. Let $HU(s)$ = the hedonic utility that accrues to oneself and $HU(o)$ = the hedonic utility that accrues to others.

<i>act</i>	<i>HU(s)</i>	<i>HU(o)</i>	<i>HU</i>	<i>moral status</i>
a_1	+10	+5	+15	permissible
a_2	+5	+10	+15	permissible or supererogatory (?)
a_3	+10	0	+10	impermissible

Lesson: Either HAU leaves no room for supererogatory acts or there is no more moral reason to perform a supererogatory act than there is to perform some permissible but non-supererogatory alternative.

15. Mr. Smith knowingly sends an innocent man to prison for ten long years. What Mr. Smith did was wrong.

Answer: Not necessarily true. It may be that punishing this innocent man maximizes hedonic utility.

Lesson: On HAU, we should punish someone if and only if doing so will maximize hedonic utility. This means that, on HAU, we should punish the innocent and we should refrain from punishing the guilty whenever doing so will maximize hedonic utility. Also, on HAU, the severity of the punishment should not necessarily be in proportion to the severity of the crime. The punishment for a less severe crime should be greater than the punishment for a more severe crime if this will maximize hedonic utility.

D. Two of HAU's Most Counterintuitive Implications

1. HAU implies that one could be morally required to murder one's own innocent daughter so as to save the life of some evil stranger even if doing so will produce only slightly more utility (.00001 hedon) than refraining from doing so would.
2. HAU implies that one could be morally required to sacrifice one's own utility, bringing one's own utility level down from -10,000 to -10,100 so as to increase some less deserving stranger's utility level from +10,000 to +10,100.00001.

E. How we might respond to HAU given its counterintuitive implications?

1. We could reject HAU, and opt for some nonconsequentialist theory in its place?
2. We could appeal to remote effects, arguing both (a) that it is only in fanciful, hypothetical cases (not real-world cases) that HAU has counterintuitive implications and (b) that we should trust our intuitions in only real-world cases.
3. We could bite the bullet, accepting that HAU has wildly counterintuitive implications even in real-world cases, but argue that this should not deter us from accepting the theory, for our moral intuitions about particular cases are not to be trusted. Utilitarians often cite three sorts of reasons for being wary of our intuitions about particular cases:
 - a. Contingency Reasons: Whereas our intuitions about what it is morally permissible to do in various particular cases seems to be contingent upon such things as our socialization, evolutionary history, and the order in which we consider various particular cases, what it is morally permissible to do in those particular cases is not contingent on such things.
 - b. Historical Reasons: As history shows, people's moral intuitions

about particular cases can be wildly mistaken. Consider, for instance, slavery.

- c. Consistency Reasons: Our moral intuitions about particular cases are often inconsistent with each other and with our other intuitions.

- 4. We could reject HAU, and opt for a different sort of consequentialist theory in its place.

II. Traditional Act Consequentialism (TAC)

A. Defined

Traditional Act Consequentialism: An act is morally permissible if and only if it maximizes the (non-moral) good, and if an act is morally permissible, then what makes it so is ultimately the fact that it maximizes the good.

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

Compare the above definition to this one: An act *maximize the good* if and only if it produces at least as much good as any available alternative act would. How does this differ from the above?

Imagine that there are three acts available to an agent: a_1 , a_2 , and a_3 . And let's suppose that these acts will produce the following three outcomes: o_1 , o_2 , and o_3 , respectively. Further suppose that although both o_2 and o_3 are better than o_1 , o_2 and o_3 are incomparable, meaning that it is not the case that they are equally good nor is it the case that one is better than the other. On the former definition, it follows that a_2 and a_3 both maximize the good, but, on the latter definition, it follows that none of the available act alternatives maximizes the good. Thus if we want TAC to yield determinate moral verdicts, we should prefer the former definition of 'maximizes the good'.

B. How TAC Differs from HAU

HAU is a species of act utilitarianism, which is itself a species of act consequentialism. Act consequentialism differs from act utilitarianism in that it is not committed to any particular theory of value, whereas act utilitarianism is: specifically, to welfarism—the view that welfare is the only intrinsic good and that ill-fare is the only intrinsic evil. And act utilitarianism differs from hedonistic act utilitarianism in that it is not committed to any particular theory of welfare, whereas hedonistic act utilitarianism is: specifically, to hedonism—

the view that pleasure is the only thing that enhances a person's welfare and that pain is the only thing that diminishes a person's welfare.

C. Traditional Act Consequentialism and the Assessment of Outcomes

1. *An Act's Outcome*: As Shaw notes: "When consequentialists refer to the results or consequences of an action, they have in mind the entire upshot of the action, that is, its overall outcome. They are concerned with whether, and to what extent, the world is better or worse because the agent elected a given course of conduct. Thus, consequentialists take into account whatever value, if any, the action has in itself, not merely the value of its subsequent effects" ("The Consequentialist Perspective," p. 6). Thus, an act's outcome includes much more than just its causal consequences.
2. *An Agent-Neutral Ranking of Outcomes*: On TAC, every agent will rank the same set of outcomes identically. If one outcome is better than another, then it is simply better (that is, better *tout court*)—not merely *better for* the agent or *better from* the agent's point of view.
3. *Assessing Outcomes in Terms of Non-Moral Goodness*:

On TAC, the goodness or badness of an act's outcome is to be assessed completely independently of any assessment of whether its act or any of its alternatives are morally permissible or impermissible. Moreover, on TAC, the goodness or badness of an outcome is to be assessed independently of any moral assessments at all: including assessments as to what distributions would be just, what people morally deserve, what agents morally ought to desire, what moral rights people have, etc. That is, outcomes are, on TAC, to be ranked solely in terms of their non-moral goodness.

4. *Assessing Outcomes in Terms of Actual Value, not Expected Value*:

On TAC, the value of an act's outcome is its actual value, not its expected value. Suppose that one places a bet with the devil, risking one hundred hedons on a coin toss such that one will receive an additional two hundred hedons in life if the coin lands heads but lose one hundred hedons if the coin lands tails. The actual value of this bet is +200 hedons if the coin lands heads and -100 hedons if the coin lands tails. The expected value, by contrast, is $(.5 \times 200) + (.5 \times -100) = 50$ hedons, and the expected value is the same whether the coin lands heads or tails. Further suppose that one does place the bet and that the coin lands heads. And suppose that the only alternative to placing the bet was refraining from doing so, which has both

an actual and expected value of +100 hedons. On objective consequentialism, placing the bet was morally permissible, for its actual value was greater than that of its alternative: that is, $200 > 100$. On subjective consequentialism, placing the bet was morally impermissible, for its expected value was less than that of its alternative: that is, $50 < 100$.

D. Its Appeal

“One thing...all [consequentialist theories] share...is a very simple and seductive idea: namely, that so far as morality is concerned, what people ought to do is to minimize evil and maximize good.... On the face of it, this idea, which lies at the heart of consequentialism, seems hard to resist. For given only the innocent-sounding assumption that good is morally preferable to evil, it seems to embody the principle that we should maximize the desirable and minimize the undesirable, and that principle seems to be one of the main elements of our conception of practical rationality....” —Samuel Scheffler, *Consequentialism and Its Critics*, pp. 1-2.

E. The Self-Effacing Nature of Traditional Act Consequentialism

As Shaw notes, “consequentialists of all stripes agree that to promote the good effectively, we should, at least sometimes, rely and encourage others to rely on secondary rules, precepts, and guidelines. Moreover, it is widely agreed among consequentialists that the full benefit of secondary rules can only be reaped when they are treated as moral rules and not merely as rules of thumb or practical aids to decision making. Having people strongly inclined to act in certain rule-designated ways, to feel guilty about failing to do so, and to use those rules to assess the conduct of others can have enormous utility. This is because it produces good results to have people strongly disposed to act in certain predictable ways, ways that generally (but perhaps not always) maximize expected benefit.” (“The Consequentialist Perspective,” pp. 14-15).

F. Rejecting TAC

Many of those who reject TAC believe that acting so as to maximize the good is either sometimes morally impermissible or sometimes morally optional, or both. In other words, many of those who reject TAC believe that there are either agent-centered constraints or agent-centered options, or both.

It is also possible to reject TAC and accept that agents are always morally required to maximize the good. One could do so if one denied that acts are permissible in virtue of maximizing the good, holding instead that acts

maximize the good in virtue of being permissible. That is, one could reject TAC by accepting what I'll call Foot's Thesis:

FT Wx is better than any available alternative world only if S is morally required to perform x , where Wx is the possible world that would be actualized were S to perform x .¹

To sum up, there are three ways to reject TAC:

- i. Hold that it is sometimes permissible to fail to maximize the good.
- ii. Hold that it is sometime impermissible to maximize the good.
- iii. Hold that an act maximizes the good only if, and in virtue of the fact that, it is morally permissible.

1. Agent-Centered Constraints

Agent-centered constraints limit what agents are allowed to do, even in the pursuit of the best consequences. They limit an agent's freedom to permissibly pursue the best consequences, impersonally construed, and they include both restrictions and special obligations:

[Agent-centered constraints go by various other names, including: *deontological constraints*, *agent-relative constraints*, *agent-centered restrictions* (Scheffler), and *side constraints* (Nozick).]

a. Restrictions, the Too Permissive Objection, and the Paradox of Deontology

Agent-centered restrictions prohibit agents from performing certain types of acts (e.g., murder) even for the sake of maximizing the good—even for the sake of preventing numerous others from performing comparable instances of that same act-type.

[Agent-centered restrictions go by various other names, including: *impersonal constraints* (Brink).]

¹ The following stronger view is incoherent: Wx is better than any available alternative world *if and only if* S would be morally required to perform x . To see this, suppose that I have the choice of pushing or not pushing button A, and assume that my pushing button A will prevent you from pushing button B, where your pushing button B would kill one but prevent five other killings. If I'm morally required to perform $\sim A$, then this stronger view would imply that the world in which I perform $\sim A$ and you perform B is better than the world in which I perform A and you perform $\sim B$. But if you're morally required to perform $\sim B$, then this stronger view would also imply that the world in which I perform A and you perform $\sim B$ is better than the world in which I perform $\sim A$ and you perform B. Thus this stronger view implies a contradiction.

Because TAC denies that there are any agent-centered restrictions, it is often said to be too permissive, permitting us to maim, murder, and massacre innocent people if doing so will maximize the good. This is known as the *too permissive objection*.

In reply, traditional act consequentialists often argue that there is something paradoxical about agent-centered restrictions. Suppose that your violating one such restriction (call it R) would result in their being fewer violations of R overall. According to the proponent of agent-centered restrictions, your violating R would, nevertheless, be morally impermissible. This can seem puzzling and, perhaps, even paradoxical. "If the non-violation of R is so important, shouldn't that be the goal? How can a concern for the non-violation of R lead to the refusal to violate R when this would prevent more extensive violations of R?" (Nozick 1974: 30 [slightly modified]) (Shaw, "The Consequentialist Perspective," p. 18). This is known as the *paradox of deontology*, because deontologists endorse agent-centered restrictions.

b. Special Obligations

According to some nonconsequentialist theories, agents have certain *special obligations* that are specific to them as individuals given their particular relationships and history. These include obligations arising out of past acts (e.g., the obligation to keep one's promises) and also those obligations that come with occupying certain roles (e.g., professional and familial obligations). It is wrong to violate such obligations even if doing so would maximize the overall good.

[Special obligations go by various other names, including: *personal constraints* (Brink) and *associative duties* (Brink).]

2. Agent-Centered Options and the Too Demanding Objection

An *agent-centered option* is an option to either safeguard one's personal interests or to sacrifice some of those interests for the sake of doing more to promote the good. These options provide agents with the freedom to permissibly act to further their own interests out of proportion to their weight in the impersonal calculus.

[Agent-centered options go by various other names, including: *agent-relative options* and *agent-centered prerogatives* (Scheffler).]

Since TAC denies that there are any agent-centered options, it requires us always to sacrifice our own interests if we can thereby do more to promote the good. Since most of the time we can do more to promote the good by sacrificing our own interests, TAC is extremely demanding. This is known as the *too demanding objection*.

III. Departing from Traditional Act Consequentialism

A. Two Ways of Departing from Traditional Act Consequentialism

There are two ways of departing from traditional act consequentialism while still remaining consequentialist: (1) departing from *act* consequentialism, adopting *rule* consequentialism instead and (2) departing from *traditional* forms of act consequentialism, adopting *nontraditional* act consequentialism instead.

B. Departing from Act Consequentialism: Rule Consequentialism (RC)

Rule consequentialism holds that the permissibility of an act depends not on the goodness of its consequences, but on whether or not it is in accordance with a certain code of rules, selected for its good consequences.

Rule Consequentialism: An act is morally permissible if and only if it is permitted by the code of rules whose internalization by the overwhelming majority of agents would maximize the good, and if an act is morally permissible, then what makes it so is ultimately the fact that it is permitted by the code of rules whose internalization by the overwhelming majority of agents would maximize the good (Hooker, *Ideal Code, Real World*, p. 32).

[Call the clause that begins “and if an act is morally permissible, then what makes it so is...” the *determination clause*. From here on out, I will for the sake of brevity omit the determination clause when I state a moral theory, but you should take it to be implicit, as some determination clause is essential to any full statement of a moral theory.]

One has internalized a rule if and only if one has acquired the dispositions and character traits that dispose one to act as the rule instructs, regardless of whether or not there is the threat of external sanction for the failure to do so.

C. Departing from Traditional Forms of Act Consequentialism: Nontraditional Act Consequentialism (NAC)

1. The basic idea:

Nontraditional Act Consequentialism: An act has the moral status that it does solely in virtue of what reasons there are for and against preferring its outcome to those of its available alternatives, such that, if S is morally required to perform x , then, of all the outcomes that S could bring about, S has most reason to want Wx to obtain.

2. Its appeal:

Whenever we face a choice of how to act, we also face a choice of which of various possible worlds to actualize. It is, after all, through our actions that we affect the way the world goes. Since rational action must aim at affecting the way the world goes—at producing some result, the first question that we must ask ourselves as moral agents is: “Which of the possible worlds that I am able to actualize do I have most moral reason to want to be actualized?” This is an important question, for it seems that I must know what result I morally ought to aim to produce before I can know what act (i.e., what means) I morally ought to perform. Thus, on NAC, the compelling thought is that what an agent morally ought to do just depends on which of the possible worlds that she can actualize is the one that she ought to prefer to all the rest, such that if S is morally required to perform x , then, of all the outcomes that S could bring about, S has most reason to want Wx to obtain.

Scheffler holds that the compelling idea that lies at the heart of TAC is the thought that “we should maximize the desirable and minimize the undesirable.” Now to say that some outcome is desirable is to say that it is an outcome that everyone ought, other things being equal, to desire. But what if there are certain outcomes that only the agent has reason to desire or that the agent has more reason to desire than others do? For instance, even if it is desirable for there to be as few murders as possible, it might be that I morally ought to prefer the state of affairs in which *two others* commit one murder each to the state of affairs where *I* commit one murder. What should I do if this is the case? Well, the compelling idea that lies at the heart of NAC is thought that each individual should act so as to bring about the outcome that she has most reason to desire. So if I should prefer the state of affairs in which two others commit one murder each to the state of affairs where I commit one murder, then I should refrain from committing murder so as to prevent these two others from committing one murder each.

3. Some Forms of Nontraditional Forms of Act Consequentialism

a) Departing from Agent-Neutrality: Agent-Relative Act Consequentialism

Agent-Relative Act Consequentialism: S's performing x is morally permissible if and only if it will actualize the possible world that, of all the possible worlds that S might actualize, is the one that S has most reason to want to be actualized.

One version:

Ethical Egoism: S's performing x is morally permissible if and only if S's doing so would maximize S's utility.

b) Departing from Maximization: Satisficing Act Consequentialism

Satisficing Act Consequentialism: An act is morally permissible if and only if its consequences are good enough.

Two versions:

Absolute Level Satisficing Utilitarianism: There is a number, n , such that: An act is morally permissible if and only if either (i) it has a utility of at least n or (ii) it maximizes utility (Bradley, "Against Satisficing Consequentialism," p. 98).

Comparative Satisficing Utilitarianism: There is a fraction, n ($0 < n < 1$), such that: An act is morally permissible if and only if its utility plus [(the utility of a maximizing alternative – the utility of a minimizing alternative) multiplied by n], is at least as great as the utility of a maximizing alternative (Bradley, "Against Satisficing Consequentialism," p. 98).

c) Departing from a Criterion of Rightness: Scalar Act Consequentialism

Scalar Act Consequentialism: Of any two acts that S might perform at a given moment, one act is morally better than another to the extent that its outcome is better than the other's outcome.

d) Departing from Non-Moral Goodness: Justicized Act Utilitarianism

Justicized Act Utilitarianism: An act is morally permissible if and only if no available act alternative would have a lower universal injustice level than it would (Feldman, *Confrontations with the Reaper*, p. 185).

The universal injustice level is the sum of the individual injustice levels of each person.

The individual injustice level of a person is the amount of utility that she has that is above or below the amount of utility that she morally deserves.

e) Departing from Actual Value: Subjective Act Consequentialism

Subjective Act Consequentialism: An act is morally permissible if and only if it maximizes expected value.

f) Departing from a Single Ranking of Outcomes: Dual-Ranking Act Consequentialism

Dual-Ranking Act Consequentialism: S's performing x is morally permissible if and only if there is no available act alternative that would actualize a possible world that S has both (i) more moral reason to want to be actualized than to want Wx to be actualized and (ii) at least as much reason, all things considered, to want to be actualized as to want Wx to be actualized.

Let Wx be the possible world that would be actualized were S to perform x .

One version:

Self/Other Utilitarianism: S's performing x is morally permissible if and only if there is no available act alternative that would produce both (i) more utility for others (i.e., those other than S) than x would and (ii) at least as much overall utility as x would (Sider, "Asymmetry and Self-Sacrifice," p. 128).

Let $U(s)$ = the utility that accrues to S, $U(o)$ = the utility that accrues to others, U = the overall utility.

<i>act</i>	<i>U(s)</i>	<i>U(o)</i>	<i>U</i>	<i>moral status</i>
a ₁	+70	+10	+80	merely permissible
a ₂	+20	+5	+25	impermissible
a ₃	+10	+15	+25	supererogatory
a ₄	-10	+20	+10	supererogatory and morally best

g) There are still other ways of departing from traditional act consequentialism (e.g., world consequentialism, multiple-act consequentialism, person-based consequentialism, etc.), but I'll save our discussion of these other nontraditional forms of consequentialism for a later day.