

in *Social Science Encyclopedia* (Routledge, forthcoming, 2006).

Consequentialism

Ethics in Practice, 3rd edition, edited by Hugh LaFollette

(Blackwell Publishers, forthcoming, 2006)

Peter Vallentyne, University of Missouri-Columbia

1. Introduction

The three main general approaches to moral theory are consequentialism, deontology, and virtue theory. I shall describe and assess consequentialism. First, however, I shall make a few background remarks on morality.

Morality is normative. It is concerned with how the world should be—as opposed to how it is. More specifically, it is normative in the sense of being concerned with what is permissible (right, acceptable) and with what is good (desirable) and bad (undesirable). There are, however, many normative perspectives. There is, for example, permissibility and goodness from the perspective of rational self-interest, legal permissibility, and aesthetic goodness. Throughout, unless otherwise specified, permissibility and goodness should be understood as moral permissibility and goodness.

Morality can assess many different kinds of things: actions, states of affairs, character-traits, social institutions, policies, and so on. To simplify our task, we will focus on the moral assessment of *actions*, which is arguably the most central form of moral assessment. Moreover, we shall focus primarily on the moral *permissibility* (as opposed to goodness) of actions. An action is permissible if and only if it is acceptable (not wrong) to perform it.¹ Our central question, then, is what determines whether an action is morally permissible?

Act consequentialism holds that the permissibility of an action is determined by how good its consequences are. For example, if shooting an innocent person would have very good overall consequences compared to its alternatives (e.g., because it saves millions of lives), then it is permissible to do so, but it is not permissible to do so if it would have relatively very bad consequences. We shall explore this and related ideas.

2. Act Utilitarianism

Act consequentialism holds that the permissibility of an action is based on how good its consequences are compared with those of its feasible alternatives. The idea is that, in a given choice situation, an agent has some number of actions that she can perform. These actions are the set of feasible alternatives and each alternative action has different consequences. The consequences include *everything* that will happen *in the world* if the action is performed. These include events in the distant future and in distant places. Suppose, for example, that, if I help a young child with her homework, she will go on to become President of the United States, avert nuclear war, and thereby ensure that a certain island still exists 1000 years from now. These are all part of the consequences of my action. The effects on everything everywhere are part of its consequences. One qualification is in order: Given that the effects of actions are typically probabilistic, the consequences of an action must be understood as specifying the probability of various states of affairs (10% chance of peace and 90% chance of war if I do this). For simplicity, however, we will typically consider simple cases where the actions produce their effects with certainty.

The historically best-known act consequentialist theory is *act utilitarianism*. It arose

primarily in Great Britain during the 1600s and 1700s, when time social thinkers were beginning to challenge the traditional social, economic, and political systems (e.g., monarchies) and their justifications in terms of God’s commands. The utilitarian emphasis was on designing and justifying social structures in terms of promoting human *wellbeing*, where a person’s *wellbeing* is a matter of how well her life goes for her overall. (Utilitarians use “utility” as a synonym for wellbeing, but we’ll generally use the latter term.) The most famous proponents are Jeremy Bentham (1748-1832), and John Stuart Mill (1806-1873).²

Act utilitarianism consists of the following two claims:

Maximizing Act Consequentialism: An action is permissible if and only if its consequences are morally maximally good (i.e., at least as good as those of its feasible alternatives).

Utilitarian Value: One state of affairs is morally at least as good as another if and only if the total individual wellbeing it includes is at least as great as the total individual wellbeing included in the other.³

We can clarify the nature of these claims by considering the following example, where there are just three feasible actions and just three people in the world:

	<u>Consequences for Wellbeing</u>			
	Jane	Mary	John	Total
Action 1	40	30	20	90
Action 2	30	20	40	90

Action 3 0 10 20 30

Example 1

Here, for example, Action 1 produces 40 units of wellbeing (e.g., happiness) for Jane, 30 for Mary, and 20 units for John. Utilitarian Value assumes that individual wellbeing can be measured on such a quantitative scale (which is controversial, and will be discussed below). Moreover, it tells us that the moral value of actions is determined by the *total* wellbeing, and thus we can consider only the total column in this example. Maximizing Act Consequentialism then tells us that Action 3 is impermissible (wrong) because it is less good than Action 1 (and also less good than Action 2). Action 1 and Action 2 are each judged permissible, because each produces as much wellbeing as possible (90). Because each is permissible, neither is obligatory (required). Each is optional (permitted but not required). The agent is required to choose either Action 1 or Action 2, but she is morally free to decide which.

We shall assess Maximizing Act Consequentialism below. First, however, we shall assess Utilitarian Value. One attractive feature of this view is that it holds that the moral goodness of states of affairs depends on how well people's lives go (i.e., their wellbeing). More specifically, it holds that making everyone better off makes things morally better. A second attractive feature is that everyone's wellbeing is given equal consideration. All individuals—rich or poor, man or woman, black or white—are given the same weight. In the above example, for instance, the wellbeing of Jane, Mary, and John counts equally, no matter what their race, religion, sex, etc.⁴

It's important to note that Utilitarian Value leaves open what determines a person's wellbeing. Early utilitarians (e.g., Bentham) took wellbeing to be something like the quantitative net balance of pleasure over pain, but this takes an unduly narrow view of the types of joy and

sorrow that are relevant for wellbeing. Later utilitarians (e.g., Mill) distinguished between higher and lower pleasures and emphasized that the quality of the pleasure and pains must also be factored in. For example, because humans typically have more sophisticated cognitive faculties, they typically are capable of much higher quality pleasures and pains. Thus, in Mill's famous words (in Ch. 2 of *Utilitarianism*), "It is [typically] better to be a human being dissatisfied than a pig satisfied". Contemporary utilitarians tend to take a still broader view of wellbeing. They claim that happiness matters a lot, but it is not the only thing that matters for wellbeing. For most people, decreasing their happiness just slightly and significantly increasing their accomplishments and the quality of their relationships with others makes their lives go better. Consider, for example, a man who is very happy, but doesn't know that his wife, whom he adores, doesn't love him and cheats on him regularly. Wouldn't his life be better if his wife loved him and didn't cheat on him, even if his happiness (for other reasons) were slightly lower? Wellbeing, on this account, is not *purely subjective* in the sense that it does not depend solely on the individual's state of mind.

The issues are complex, and I'll here simply mention two of the main accounts of wellbeing that have been developed. *Preference Utilitarianism* claims that wellbeing is a matter of how well a person's informed self-regarding preferences are satisfied. This is partly subjective (the preferences) and partly objective (how well the world satisfies those preferences). The husband in the above case would probably prefer to be slightly less happy and have his wife love him more and be faithful to him. *Perfectionistic (or Ideal) Utilitarianism* claims that there are some things that are objectively good for a person even if she doesn't care about them (in her preferences). For example, one might claim that increasing a person's accomplishments or

positive intimate relationships makes her life go better even if she doesn't care about them. In what follows, we'll leave open what the correct account of wellbeing is.

Utilitarianism Value faces several important objections. One is that it presupposes that well-being can be *quantified very precisely*. In the above example, for instance, it does not merely assume (as is plausible) that there are facts about when Jane has more wellbeing (e.g., that Action 1 gives her more wellbeing than Action 2). Utilitarian Value also presupposes that there are facts about *how much more* wellbeing Jane has (e.g., that the difference in wellbeing for Jane between Action 1 (40 units) and Action 2 (30 units) is less than the difference in wellbeing for her between Action 2 (30 units) and Action 3 (0 units). Moreover, it further assumes that wellbeing is *interpersonally comparable* in the sense that there are facts about how one person's increase in wellbeing compares with that of another. Utilitarianism Value presupposes, for instance, that, in moving from Action 2 to Action 1, Jane gains less wellbeing (10 units) than John loses (20 units). It's not clear, however, that such precise quantitative measurement of wellbeing is possible even in principle. Of course, the possibility of such measurement depends on how exactly wellbeing is understood. Defenders of Utilitarian Value must thus supply a clear account of wellbeing that is quantitatively measurable and interpersonally comparable. If no such account can be given, then Utilitarian Value rests on a false presupposition.

A second objection to Utilitarian Value is that, because it focuses solely on *total* wellbeing, it is *insensitive to how wellbeing is distributed among individuals*. To see this, consider the following example:

Consequences for Wellbeing

Jane Mary John Total

Action 1	0	0	99	99
Action 2	33	33	33	99
Action 3	50	10	20	80

Example 2

Here, for example, Action 1 produces 0 units of wellbeing for Jane and Mary and 99 units of wellbeing for John. Utilitarian Value tells us that the moral value of these actions is determined by the total wellbeing, and thus we can consider only the total column. Action 1 and Action 2 are equally good (with a total of 99 each) and each is morally better than Action 3. Let us here grant that Action 3 is worse than the other two. Is it true that the first two actions are equally good? They have the same total, but wellbeing is more equally distributed in Action 2. If no one has any special claim to more wellbeing (e.g., because of working harder), the more equal distribution, it seems, is better. Utilitarian Value, however, leaves no room for distributive considerations such as equality or priority for those who are worse off. Given that some kind of distributive consideration seems relevant, this is a very significant objection to utilitarianism.

Utilitarians have a partial reply to the charge that they give no priority to the worse off and no consideration to equality. They appeal to *decreasing marginal wellbeing (utility)* from resources (or money). The idea is that, for a given person, the increase in wellbeing that results from having an additional unit of the resource (e.g., additional apple or additional dollar) decreases the more resources that individual has. Thus, for example, if a person has no money, acquiring a dollar gives a significant benefit in wellbeing (e.g., makes the difference between starving and not starving), but, if that person already has a million dollars, acquiring an additional dollar has a trivial impact on wellbeing. Given that individuals have roughly the same

dispositions for decreasing marginal wellbeing, this means that, all else being equal, utilitarians will favor (because it will increase total wellbeing) shifting resources from the rich to the poor (and thus favor equality). Of course, there are other factors, such as incentive effects, that must be factored in as well, and they may push in the opposite direction. (For example, providing aid to the poor may reduce their incentive to work, and their reduction in work may reduce the total wellbeing.) Still, in light of the roughly equal decreasing marginal wellbeing of all, utilitarianism has some significant tendency to favor more equal distributions of resources over less equal ones.

The important point to note, however, is that equality of resources (e.g., money) is not equality of wellbeing. A person who regularly suffers severe depression may have the same amount of money (and other resources) as someone with a joyful disposition, but their levels of wellbeing will be very different. Thus, Utilitarianism's sensitivity to the distribution of resources does not automatically make it suitably sensitive to the distribution of wellbeing.

A third and related objection to Utilitarian Value is that it is *insensitive to what choices individuals make*. It doesn't care, for example, what good or bad deeds an individual performs. Everyone's wellbeing counts equally. Thus, for example, it views 100 units of wellbeing for a do-gooder and 0 units for an evildoer as equally as good as 0 units for the do-gooder and 100 for the evildoer (since they have the same total). There is no room for some individuals *deserving* more than others because of their choices.

In light of the last two objections, many consequentialists have abandoned Utilitarianism Value. Instead, they appeal to some other theory of moral value (of what makes one state of affairs morally better than another). Their preferred theory of value might be sensitive equality of wellbeing, how well wellbeing matches desert, or even appeal to considerations that have

nothing to do with wellbeing (e.g., the promotion of knowledge and beauty). In abandoning Utilitarian Value, such consequentialists are abandoning utilitarianism, but they are not abandoning the act consequentialist idea that the permissibility of actions is determined by how good their consequences are. In what follows, then, we shall assume that some suitable theory of value has been adopted, but leave open its exact content. We shall thus focus on consequentialism in general (as opposed to the utilitarian version thereof).

3. Act Consequentialism.

Act Utilitarianism, we saw above, consists of Utilitarian Value plus the following thesis:

Maximizing Act Consequentialism: An action is permissible if and only if its consequences are morally maximally good (i.e., at least as good as those of its feasible alternatives).⁵

We shall now focus on this principle (which leaves open how moral goodness is determined).

Maximizing Act Consequentialism has two main attractive features. First, it is hardheaded and focuses on what the consequences of actions will be. It rightly holds that the permissibility of an action depends on how morally good its effects would be. Second, Maximizing Act Consequentialism requires agents to do the best they can. It requires them to perform the morally best action (in terms of consequences) that they can. Something about that seems right.

Maximizing Act Consequentialism faces, however, several important objections. One is that it is impossible, or at least counterproductive, to calculate the consequences of one's actions each time one makes a choice. To do this adequately, after all, one would have to have

knowledge of all the consequences of one's actions for the entire world for all of time. No one has anything close to this knowledge, and thus it is practically impossible to apply a Maximizing Act Consequentialist principle with any reliability. This objection, however, is easily met. First, note that Maximizing Act Consequentialism is not a *decision-procedure* that agents are supposed to follow consciously when making choices. It is instead a *criterion of permissibility*. It specifies the conditions that determine whether an action is permissible. Thus, for example, an agent might steal someone's yacht with the sole aim of getting rich by selling it. If the yacht belongs to terrorists who were going to use it to kill thousands of people, this action might have the morally best consequences, and thus be judged permissible. The fact that the agent was not consciously attempting to accomplish the moral good is not deemed relevant.

This reply, however, does not get Maximizing Act Consequentialism completely off the hook. After all, agents should at least sometimes consciously reflect upon what will have the best consequences, and given their limited time and information, this seems practically impossible or counterproductive. Given all the needed information, for example, agents could spend all their time gathering and processing information and never make a substantive choice. Maximizing Act Consequentialists have, however, an answer to this problem. It is that we should adopt various rules of thumb (e.g., don't kill, don't steal, don't lie) the application of which generally produce the best results. Under normal circumstances, there is no need to perform consequentialist calculations. Indeed, to do so will normally be wrong (since it will waste time and not produce the best consequences). Instead, one typically just applies the rules of thumb. It is only in special circumstances—in which one has special reason to think that one's choice will have unusual or unusually important consequences—that one should perform consequentialist calculations. Of course, agents will make lots of mistakes (i.e., perform impermissible actions), but this is just a

fact of life. Morality, consequentialists claim, is very complex, and human beings have only limited time and knowledge. It is therefore not surprising that we often make mistakes. We should simply do the best we can.

This reply, I believe, adequately meets the objection. The next objection, however, is more powerful. It claims that Maximizing Act Consequentialism is *too demanding* in that it typically judges permissible only a very small percentage of an agent's feasible options. Indeed, if there is just one action that has the best consequences, there will be just one action that is permissible. Assuming (as we shall) that ties in moral value are relatively rare, then there will typically only a few permissible actions in any given choice situation. There are two aspects of the demandingness objection. One is that, under most realistic conditions, it typically requires agents to make significant sacrifices of their own wellbeing to maximize moral goodness. The objection here is *not* that Maximizing Act Consequentialism *sometimes* requires agent's to make significant sacrifices; all plausible moral theories have this feature. Any plausible theory, for example, will typically judge it impermissible to steal a million dollars, even though one can get away with it and would greatly benefit from the result. The objection here is that Maximizing Act Consequentialism *frequently* requires significant sacrifices from agents. It holds that typically it is wrong to spend money (e.g., for restaurants, clothes, or CDs) or time (e.g., watching TV, talking with friends) for one's own enjoyment, since morally better consequences can be obtained by using this money or time in other ways (e.g., helping the needy). Of course, such activities are not always wrong, since the most effective way of promoting moral goodness typically involves occasionally pampering oneself (e.g., to recharge one's batteries). Most of the time, however, Maximizing Act Consequentialism judges it impermissible to devote more than minimal time or resources to oneself.

A second aspect of the demandingness objection is that Maximizing Act Consequentialism leaves no room for agent's to favor their loved ones or others with whom they have special relationships. If one can save either one's own child (or friend) or a stranger, one is required to produce the morally best consequences. There is no room (except in rare cases of ties in goodness) to choose among various permissible options and favor one's loved ones.

In light of the demandingness objection, some consequentialists have rejected Maximizing Act Consequentialism in favor of:

Satisficing Act Consequentialism: An action is permissible if and only if its consequences are morally adequate.

The idea is that morality requires that the consequences be adequate but not that they be maximally good. There are different kinds of criteria for adequacy. Consequences might be judged adequate if they are better than those of at least 50% of the alternatives, or if they do not make things worse than doing nothing. We shall not explore the important task of developing plausible criteria of adequacy. The important point is that, as long as the criterion requires significantly less than the best consequences, there will be significantly more moral freedom for agents to promote their own wellbeing or that of their loved ones. As long as the criterion of adequacy is sufficiently weak, the objection of demandingness can be avoided.

Let us therefore focus on the following broader class of theories:

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

Maximizing Act Consequentialism and Satisficing Act Consequentialism are each forms of Act Consequentialism. The former hold that only the best possible consequences are good enough, whereas the latter hold that adequate consequences are good enough.

Act Consequentialism is subject to the objection that it gives little protection to individuals from gross interference from others. It judges it permissible to kill, torture, lie to, and steal from innocent individuals, whenever this produces sufficiently good outcomes. All that matters is the overall result. If it is good enough, then it is permissible to do whatever it takes to produce that result. The ends can justify any means. Nothing is ruled out in principle.⁶ Most of us recoil at this idea. Even assuming that there is some kind of duty to promote good consequences, surely there are limits on the permissible means of how we treat others.

Of course, things are not quite so straightforward. Typically, treating people in horrible ways (killing, torturing, etc.) has very bad consequences. For example, it will cause people in general to fear being a victim of abuse. Thus, as a rule of thumb, Act Consequentialism will not favor such treatment. It will be only under special circumstances that act consequentialism will judge such treatment permissible. If circumstances are truly extraordinary (e.g., it is necessary to kill an innocent person to save 100,000 innocent people), then perhaps it is not wrong to do so.

Nonetheless, Act Consequentialism faces a problem here. Let us grant for the sake of argument that, under truly extraordinary circumstances, it is permissible to treat innocents in horrible ways and that, as a rule of thumb, Act Consequentialists will not favor gross mistreatment of individuals. The important point is that, all else being equal, Act Consequentialism will, for example, judge it permissible to torture and kill *one* innocent person when this is the only way to avoid *two* innocent people from being comparably tortured and killed by others. Of course, in

practice Act Consequentialists will apply rules of thumb and be cautious about torturing and killing innocents. The point here is that the circumstances under which Act Consequentialism judges it permissible to engage in such abuses is not limited to rare and extraordinary cases of avoiding social catastrophe. Most of us doubt that the ends justify the means as frequently as Act Consequentialism says they do.

4. Constrained Act Consequentialism and Rule Consequentialism

In response to the above criticism about the ends not always justifying the means, one might consider adopting the following modification:

Constrained Act Consequentialism: An action is permissible if and only if, of those actions that satisfy certain specified constraints, it has consequences that are sufficiently morally good.

Here we suppose that there are some independently supplied constraints that actions must respect. They might rule out, for example, killing and harming innocents, lying, breaking promises and agreements, and theft. These constraints might be based on the rights of individuals or they might have some other source.

With an appropriate set of constraints, Constrained Act Consequentialism avoids the problem of readily sacrificing individuals for the greater moral good. Moreover, if it takes a non-maximizing form, it avoids the problem of being excessively demanding of agents. More generally, it can capture at least the broad outlines of common-sense morality. We have some obligation to promote moral goodness, but we are left a fair amount of liberty in how we do so. We are not frequently required to make major sacrifices of our own wellbeing and are typically

permitted to give special treatment to our friends and loved ones. Moreover, the ends do not always justify the means: there are some constraints on the permissible ways of promoting the moral good.

The important point to note is that Constrained Act Consequentialism is not a form of Act Consequentialism. It is rather a mixed theory: It has a deontological component (the constraints) and a consequentialist component (the duty to promote the good). To see that it is not a form of Act Consequentialism, it suffices to note that it sometimes judges impermissible the action with the best consequences. Suppose, for example, that all else is equal, and that, by killing one innocent person, you could save ten innocent people from being killed by a terrorist. We may suppose here that it is morally better that one innocent person be killed than that ten innocent people (with comparable lives) are killed. Thus, killing one innocent person has the best outcome, but Constrained Act Consequentialism judges it impermissible. This is, of course, just an example of the fact that it rejects the view that the ends always justify the means.

If one accepts Constrained Act Consequentialism, then one is a *deontic pluralist* in that one holds that there is more than one fundamental moral consideration for determining what is permissible. One such consideration is how well an action promotes the moral value of the consequences, but another is whether the action respects certain deontological constraints. Many people find such pluralism plausible, but Act Consequentialists do not. Although they can allow that there may be many considerations that are relevant for moral goodness, Act Consequentialists insist that the only moral consideration relevant for moral permissibility is how well moral goodness is promoted. Thus, the belief that there are some deontological constraints on how the good may be promoted has led some to reject Act Consequentialism.

So far, we have focused on *Act-Consequentialism*. It grounds the permissibility of actions

on their consequences. We shall now consider a form of consequentialism that appeals to the consequences of adopting (normative) *rules* (as opposed to actions). Such rules consist of practical do's and don't's, such as "Never tell a lie.", "Keep your promises", and more complex variations. Consider, then:

Rule Maximizing Consequentialism: An action is permissible if and only if it conforms to rules that, if generally followed (internalized, upheld, etc.), would have consequences that are at least as good as any feasible alternative set of rules.

Rule Maximizing Consequentialism does not assess actions on the basis of the value of their consequences. Instead, it assesses them on the basis of their compliance with selected rules, and it selects rules on the basis of the value of the consequences of their being generally followed (internalized, upheld, etc.). Rule Maximizing Consequentialism has the potential to avoid all the objections raised above. Because it does not assume Utilitarian Value, it can be sensitive to distributive considerations of various sorts (e.g., equality or desert). Because the rules with the best consequences are likely to leave agents a reasonable amount of moral freedom, Rule Maximizing Consequentialism will typically not require excessive sacrifices from agents and will leave them some significant freedom to favor their loved ones. Finally, because the rules with the best consequences are likely to give a certain basic protection to agents against interference from others, Rule Maximizing Consequentialism will recognize various constraints on how individuals may be treated.

Does this seem too good to be true? Many philosophers believe so. The main objection to Rule Maximizing Consequentialism is that it is inadequately sensitive to the consequences of

actions. Suppose, for example, that the best set of rules prohibits lying, and you can either lie to the terrorist—which will thwart her plans, thereby saving thousands—or you can do something else (tell the truth, say nothing, be evasive, etc.)—in which case the terrorist will succeed in killing thousands. Rule Maximizing Consequentialism says that you must obey the best set of rules, and that set of rules prohibits lying. Hence, Rule Maximizing Consequentialism says that lying is impermissible in this case. Act Consequentialists think that this is crazy. It is the consequences of our actions that matter, and it seems quite wrong not to lie in this case. The fact that some set of rules has the best consequences if generally followed seems irrelevant. Rule Maximizing Consequentialism seems guilty of “rule worship”.

Of course, the rules that have the best consequences may not be as simple as in the above example. Maybe, all the prohibitions are qualified by “unless doing so would have unusually good consequences” and maybe there is some kind of overriding master requirement to do whatever has unusually good consequences. This would, of course, reduce the objection of rule worship, but it also creates the danger that Rule Maximizing Consequentialism collapses into Maximizing Act Consequentialism. Whether it does so collapse is something that has been much debated, and we shall not attempt to resolve the issue here.⁷

5. Conclusion

We have surveyed some of the main forms of consequentialism. Each has its problems. This, however, does not establish that no version of consequentialism is correct. Once one examines the issues carefully, every moral theory seems to have problems of some sort. This is because our *untutored* moral judgements—which guide our judgements about what answers are correct—are not perfectly coherent. They are often based on false assumptions, confused notions, or failure to

see the full implications of a view. The real test of a theory is how good its answers are relative to our *reflective* moral judgements, which are (roughly) the judgements that we have after we have thoroughly investigated all moral and related empirical issues. Thus, given that some of our current moral judgements may be mistaken, some of the seemingly powerful objections to consequentialism may be mistaken. If so, then some version of consequentialism might be the correct theory of morality.⁸

¹ Some permissible actions are optional (permissible to perform but also permissible not perform) and some are obligatory (permissible to perform and not permissible not to perform). Most people think that, under normal circumstances, scratching one's head is morally optional but keeping a contract is morally obligatory.

² See Jeremy Bentham, *Introduction to the Principles of Morals and Legislation* (London, 1789) and John Stuart Mill, *Utilitarianism* (London: 1863). An excellent introduction to utilitarianism is J.J.C. Smart, and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973).

³ For simplicity, I here focus on Total Utilitarianism. Another version is Average Utilitarianism, according to which a state of affairs is morally at least as good as another if and only if its average wellbeing is at least as great. If the number of people is the same in both states of affairs, there is no difference between these two views. Where the number of people is different, however, the two views can diverge. The total view says, for example, that three people with 1 unit (total of 3, average of 1) is better than one person with 2, whereas the average view says the opposite.

⁴ Utilitarian Value is sometimes described as calling for “the greatest good for the greatest

number”, but this is a mistake—and indeed is incoherent. Utilitarian Value says that 99 people with 0 and one person with 100 (total 100) is better than 99 people with 1 and one person with 0 (total 99)—even though the latter is better for 99 of the 100 people.

⁵ Consequentialism is sometimes characterized more broadly so as to include moral egoism—the thesis that an action is permissible if and only if its consequences are maximally *prudentially good for the agent*. This view, however, is a non-starter as a moral theory—given that morality is concerned with everyone’s wellbeing. More generally, the most plausible versions of consequentialism are based on moral goodness, and hence I shall focus on them.

⁶ It’s worth noting that the satisficing version of act consequentialism is especially prone to this objection, since it sets the bar lower for the consequences to be good enough.

⁷ For a state of the art defense of rule consequentialism, see Brad Hooker, *Ideal Code, Real World* (Oxford: Oxford University Press, 2000).

⁸ For helpful comments, I thank Hugh LaFollette, Eric Roark, and Alan Tomhave.