



OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

ANALYSIS

Average Utilitarianisms

Author(s): T. M. Hurka

Source: *Analysis*, Vol. 42, No. 2 (Mar., 1982), pp. 65-69

Published by: [Oxford University Press](#) on behalf of [The Analysis Committee](#)

Stable URL: <http://www.jstor.org/stable/3327924>

Accessed: 16-03-2016 20:00 UTC

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Oxford University Press and The Analysis Committee are collaborating with JSTOR to digitize, preserve and extend access to *Analysis*.

<http://www.jstor.org>

AVERAGE UTILITARIANISMS

By T. M. HURKA

AN important class of utilitarian moral theories calculate value, which is what they tell us to maximize, by calculating happiness both across all the persons in human history, and across all the times in human history.¹ Classical utilitarianism is one of the theories in this class, and it uses summative principles to make both these calculations. It determines the happiness felt by a group of n persons by summing the happiness felt by all of its members; and it determines the happiness felt in a period of n years by summing the happiness felt in all of those years. Nor does it matter what order classical utilitarianism applies these two summative principles in. If it wants, it can begin by summing the happiness felt by a single person at all the times in his life; this will give it a measure of the total happiness in his life, and it can then sum these measures across lives. Or it can begin by summing the happiness felt by all the persons in the world at a single time; this will give it a measure of the total happiness in the world at that time, and it can then sum these measures across times. Whatever order classical utilitarianism applies its two principles in the result will be exactly the same; and as a result nothing turns on the order it actually adopts.

Average utilitarian theories use averaging principles at least somewhere in their calculations of happiness. But there are several such theories in the class of those which calculate across all the persons and times in human history, and in all of them a great deal turns on the order in which the principles in them are applied. Let us begin by considering the average utilitarian theories in this class which use an averaging principle to calculate happiness across persons, and a summative principle to calculate happiness across times (they are probably closest to what philosophers have had in mind when they have spoken about "average utilitarianism" in the singular). These theories can begin by summing the happiness felt by a single person at all the times in his life; this will give them a measure of the total happiness in his life, and they can average these measures across lives. Or they can begin by calculating the average amount of happiness felt by each person alive at a single moment; this will give them a measure of the average happiness per person at that moment, and they can then sum these measures across moments. But these two procedures are

¹ Utilitarians often include animal happiness in their calculations of value. I have only excluded it because this simplifies some of my arguments considerably.

not guaranteed to yield the same result. There are circumstances in which they yield different results, and we must therefore regard the theories which adopt them as different average utilitarian theories. Let us call the theory which adopts the first procedure—the one which first sums across times and then averages across persons—A1, and the theory which first averages across persons and then sums across times A2.

Both A1 and A2 have unattractive consequences, but these unattractive consequences are different. A2 requires us to kill people who will feel less than the average amount of happiness from now on into the future. Let us imagine that the world is an extremely happy place, in which almost everyone is enjoying 1000 units of happiness every year (this on a scale where the best we do today is 10 units of happiness a year). But one person is feeling only 900 units of happiness every year. Then even though this person's life is a very happy one A2 requires us to kill him. His existence at any future moment would lower the average happiness per person at that moment, and we must therefore prevent it. The fact that A2 requires us to kill below-average people in this way surely constitutes a powerful objection against it; but it is important to realize that the same objection cannot be raised against A1. A1 tells us to maximize the average happiness per human life lived; and it calculates happiness within lives using a summative principle. To end the life of a man who is enjoying 900 units of happiness every year is to decrease the amount of happiness his life will contain as a whole, and it is therefore also to decrease the contribution his life will make to the general average. Killing him will not remove his life from the list of those lived in human history, and A1 therefore has no reason to recommend it.²

Although A1 does not have unattractive consequences about killing it does have unattractive consequences in other areas. Imagine that human lives up to this point in history have been extremely happy, containing an average of 100,000 units of happiness per life lived. But from now on human lives will be slightly less happy. As a result of unfortunate climatic changes (winter has disappeared and people have to live all year long without ice and snow) human lives will contain only 90,000 units of happiness. Then A1 requires us all to stop having children, and indeed to let the human race die out. Even though the children we produced would lead very happy lives (90,000 units of happiness where the best we do today is 1000), these lives would lower the average happiness per life in human

² The claim that "average utilitarianism" (in the singular) requires the killing of below-average people has been made by a number of writers, among them James E. Meade, Richard Henson, Robert Nozick, Philip E. Devine, Peter Singer, Robert Scott, Jr., and Wayne Sumner. It is not clear, however, that these writers had A2 in mind when they made this claim, for some average utilitarian theories which calculate happiness across only *future* persons and times also require the killing of below-average people, though who counts as below-average in these theories is slightly different. Claims about the consequences of "average utilitarianism" in the singular have not always been made, it seems to me, with a full awareness of the large number of different average utilitarian theories there are.

history, and they must therefore not be lived. The fact that A1 requires us to let the human race die out in this situation surely constitutes a powerful objection against it; but it is important to realize that the same objection cannot be raised against A2. A2 calculates the average happiness per person at a single moment, but it then goes on to sum these averages across moments. To prevent the existence of people who would feel 90,000 units of happiness per person each year is to decrease the average happiness at every moment in the future from a large positive number to 0; and it is therefore also to decrease the sum of those averages across all the moments in human history.

A1 and A2 also have unattractive consequences associated with their use of summative principles across times (some of these were pointed out to me by Derek Parfit). A2, for instance, entails something like the “repugnant conclusion” associated with classical utilitarianism. Let us try to imagine an ideal (finite) course of human history, one in which the average happiness per person is very high for a very long (finite) period of time. Then if A2 is correct there is another vastly longer course of human history which is better than this one, even though the average happiness per person in this second course is never more than barely positive, and the experiences of the people in it never more than barely worth having. The second course is better because, if it is long enough, the sum of the average happiness per person at each time is greater than it is in the first course. A1 sometimes shares this unattractive consequence with A2 and sometimes does not. If the lives lived in the second course of history are also vastly longer than those lived in the first course, and therefore contain more total happiness than those in the first course, A1 joins A2 in preferring the longer second course. But if the lives in the second course are roughly the same length as those in the first course, A1 prefers the shorter first course. A1 also entails something like the repugnant conclusion. Let us try to imagine an ideal human life (of finite duration), one in which the average happiness per moment is very high for a very long (finite) period of time. Then if A1 is correct there is another vastly longer human life which it would be better to have lived than this one, even though the average happiness per moment in this second life is never more than barely positive, and the experiences it contains never more than barely worth having. A2 sometimes shares this unattractive consequence with A1 and sometimes does not. If the other lives lived alongside these two lives have an average happiness per moment which is lower than or equal to that of the second life, or if there are no other lives lived alongside the two lives, A2 joins A1 in preferring the longer second life. But if the other lives lived alongside these two have a higher average happiness per moment than the second, A2 prefers the shorter first life.

A1 and A2, then, are open to a number of objections on account of their consequences for particular moral choices. A more theoretical objection,

and one which applies to both theories equally, is suggested by a famous passage in Henry Sidgwick's *The Methods of Ethics* (7th edition, Macmillan 1907, p. 418). In this passage Sidgwick asks how an egoistic hedonist can justify giving equal weight to pleasures felt at different times in his own life and no weight to pleasures felt by different persons, when the problems raised by pleasures felt at different times and pleasures felt by different persons are from a formal point of view exactly analogous. The question Sidgwick asks about weightings can also be asked about calculations of happiness. If the problems raised by calculating happiness across persons and times are also formally analogous, as they clearly are, how can A1 and A2 justify using averaging principles to calculate happiness across persons and summative principles to calculate happiness across times? If A1 and A2 could point to some relevant difference between persons and times then of course their use of these different principles would be beyond reproach; but it is difficult to see what this relevant difference could be, and difficult to see how their use of different principles could ever be justified.

Reflection on this "symmetry objection" leads us to consider the average utilitarian theories in our class which are not open to it, because they use averaging principles both across persons and across times. There are two such theories, differing from each other in the order in which their two averaging principles are applied. The first of these theories (call it A3) begins as A1 did by calculating the happiness within a single human life, but instead of calculating the total amount of happiness in a life A3 calculates the average happiness per moment in it. Having calculated this average, however, it then follows A1 in averaging across lives, and tells us to maximize the average of the average happiness per moment in all the lives in human history. The second theory (call it A4) begins as A2 did by calculating the average happiness per person at a single moment; but instead of telling us to maximize the sum of these averages across moments it tells us to maximize their average.

Although A3 and A4 are not open to the symmetry objection, and do not entail anything like the repugnant conclusion, their consequences for other particular moral choices are even more unattractive than those of A1 and A2. A3 has the same unattractive consequences about letting the human race die out that A1 has, but it also has some unattractive consequences about killing. Imagine that a person has had an extremely happy life up to the present, enjoying 1000 units of happiness in every year that he has lived. But in the future he will enjoy only 900 units of happiness every year. Then even though these extra years would be very happy ones A3 says they must not be lived. They would lower the average happiness per moment in this person's life, and we must therefore kill him. (To see that A4 is distinct from A3 imagine that this person, who has been happier than everyone else in the past, will still be happier than everyone else in the future, even though everyone else will be happier in the future than they were in the

past. In this situation A4 says he must be kept alive.) A4 has the same unattractive consequences about killing that A2 has, but it also has some unattractive consequences about letting the human race die out. Imagine that the average happiness per person has been very high throughout human history, with every person enjoying at least 1000 units of happiness in every year of his life. But recent climatic changes mean that in the future everyone will enjoy only 900 units of happiness every year. Then even though people would be very happy in the future A4 requires us to let the human race die out. The average average happiness per person would be lowered if there were any other moments at which humans were alive, and procreation must therefore come to an end.

The last average utilitarian theories in our class for us to consider (and there are once again two of them, which we can call A5 and A6) use summative principles across persons and averaging principles across times. A5 and A6 are probably the least attractive average utilitarian theories of all in our class. They are open to the symmetry objection, and combine many of the most unattractive consequences of classical utilitarianism and A1 and A2. The identification of these consequences I will leave to the reader.

“Average utilitarianism” is usually discussed and criticized as if it were a single theory. This is a mistake. There are six different theories which can be called average utilitarian in the class of those which calculate happiness across all the persons and times in human history alone, and although these theories are all open to powerful objections the objections are in many cases importantly different.

University of Calgary

© T. M. HURKA 1982