Australasian Journal of Philosophy
Vol. 73, No. 3; September 1995

INFINITE UTILITY

James Cain

I. Introduction

If each of a pair of actions would continue to generate good or bad consequences over an infinite span of time so that both would result in an infinite net good (or both in an infinite net bad), it is difficult to see how an act utilitarian could decide which of the actions was better. Peter Vallentyne proposes what I take to be an initially plausible way that an act utilitarian might try to handle this problem.[1] He proposes the following definition and principle:

> PMU*: An action, a1 *produces more utility* than an action, a2, if and only if there is a time t such that for *any* later time t' the *cumulative* amount of utility produced by a1 up to t' is *greater* than that produced by action a2 up to t'.[2]
>
> U: An action is permissible if and only if *no alternative produces more utility*.[3]

PMU* is given as a stipulative definition of the notion of *producing more utility than*. Vallentyne states that '. . . PMU* and U seem to be necessary and sufficient for preserving the coherence and spirit of utilitarianism in cases involving infinite utility'.[4]

We shall see, however, that examples can be found which show that PMU* and U fail to give us a plausible way to extend utilitarianism to cases involving infinite utilities. We can even find cases in which, for a pair of actions, a1 and a2, up to any future point of time, t, a1 has an infinite positive utility up to t; a2 has an infinite negative utility up to t; and yet a2 is clearly the action which ought to be preferred under any reasonable extension of utilitarianism to cover cases involving infinite utility! This kind of case can arise because if we compare the respective utilities of a pair of actions on a time interval by time interval basis we may get a different result than we get by comparing utilities on an agent by agent basis.

II. First Example: The Sphere of Suffering

Imagine the following situation. We have an infinite universe in which there are infinitely many persons (and there are only finitely many persons in any given finite

[1] 'Utilitarianism and Infinite Utility', *Australasian Journal of Philosophy* 71 (1993) pp. 212-217.
[2] *Ibid.*, p. 215.
[3] *Ibid.*, p. 216.
[4] *Ibid.*, p. 217.

volume). We imagine that, with respect to some given frame of reference, the spatial locations of these people remain fixed. These beings are immortal and no other living beings exist or will come into existence. I will assume that there are no other utilities to consider beyond those of these people. Both action a1 and a2 will bring into existence a sphere one foot in diameter which will remain centred on the same point and will grow in diameter one foot per year. Action a1 brings it about that everyone within the sphere suffers disutility at a fixed finite level per unit of time, and anyone outside the sphere has a positive utility, again, at a fixed finite level per unit of time. Action a2 is similar, except that those within the sphere have a positive utility and those outside suffer. Suppose these are the only utilities to be considered. Which course of action is preferable?

Action a1 has the result that anyone who is inside the sphere from the beginning will suffer negative utility forever. Anyone else will have a finite span of positive utility after which it will become engulfed in the sphere and suffer forevermore. Action a2 has the result that everyone will at most endure a finite period of suffering before becoming engulfed in the sphere, after which they will enjoy positive utility forever. Clearly we should prefer a2 over a1. Even Bentham's method from Chapter IV of *An Introduction to the Principles of Legislation* would tell us to prefer a2, since he would have us first find the net utilities for each individual involved and then sum these up to get the total utility of the action. Action a1 would thus have infinite negative utility and a2 would have infinite positive utility.[5]

But note the following strange feature of this example. For any finite interval of time after a1 is performed, there will have existed an infinite net positive utility, for there will have been infinitely many beings outside the sphere for that whole time span and only finitely many will have been within the sphere. On the other hand, we see that, for any finite interval of time after a2, there will be an infinite negative utility, since during any time interval there will be infinitely many beings suffering outside the sphere and only finitely many will have entered the sphere. So PMU^* and U tells us that a1 is preferable. Clearly this is not a reasonable way to extend utilitarianism to cover cases of infinite utility.

III. Second Example: The Headache

In the previous example there was an infinite number of people existing at any one time and each person had an infinite life span. One might suspect that Vallentyne's account would work in cases where there are only finite populations to deal with at any given time and everyone is mortal. Our next example shows that this is not so. Imagine that we may choose between actions a3 and a4. Suppose both a3 and a4 will create a world in which one person initially exists; at the end of every month the population of the world triples, and all beings live for, say, eighty years and one month (though it would suffice for our example just to have all beings live more than two months). Action a3 will have the further result that after people have existed for one month they get a headache that lasts for the remaining eighty years of

[5] In calculating the net utilities for each individual I am ignoring Bentham's appeal to the 'propinquity or remoteness' of pleasures.

their life, but they are headache-free for their first month. Action a4 reverses this: for the first month one has a headache, but after this one is headache-free. For simplicity, we imagine that anyone with a headache has a net negative utility of one unit per month, and anyone without a headache has a net positive utility of one unit per month.

It seems clear that a4 would be the better choice, since everyone will then have a headache for one month and then be headache-free for eighty years, whereas if a3 is chosen everyone will have one month without a headache and then eighty years of pain. And a4 is the choice that would be made by a utilitarian (like Bentham) who first sums up the net utility for each individual and then calculates the overall net utility for the whole population whose utility is at stake, for a3 will have a negative net utility of -959 (= 1-(80·12)) for each person, and a4 has a positive net utility of 959 (= (80·12)-1) for each person.

Yet Vallentyne's proposal would have us choose a3, for after any finite length of time the net utility so far generated by a3 will always be ahead of the net utility so far generated by a4. This holds because the population triples at the end of each month, and so there are always at least twice as many people in their first month of life as there are living people over one month old. Thus at any point in time there will be at least twice as many people with a headache as without if a4 is performed, and there will always be at least twice as many without a headache as with a headache if a3 is performed.

IV. The Problem of Simultaneity

Though I personally have no inclination to accept utilitarianism, I have argued that the most natural way to extend utilitarianism to handle infinite utilities is not the way advocated by Vallentyne. This does not mean that Vallentyne's intuitions would be of no use for a utilitarian in treating infinite utilities. A utilitarian might want to hold the following:

> In cases where infinite utilities are involved, when choosing between actions X and Y, *if* it is the case that X maximizes utility for each individual concerned then X is the better choice. There will be a range of cases, however, in which, on the basis of considering net utilities for each individual, neither X nor Y comes out ahead. Somewhere within this range of cases Vallentyne's principle U will find its proper domain of application.

For example, suppose that we have an expanding sphere just as in the first example, except that at every instant the universe contains infinitely many mortals (everyone's life-span being exactly eighty years). As the sphere grows it encompasses more and more people, but at every instant it only contains finitely many people, whereas there are infinitely many outside the sphere. Over the whole course of the history of the universe there will be infinitely many individuals with positive net utilities and infinitely many with negative net utilities regardless of whether it is those inside the sphere that suffer (while those outside have a positive utility) or

those outside who suffer (while those inside have a positive utility). Now in a case like this, Vallentyne's account seems particularly appealing: here (at least at first glance) it does seem preferable to have the disutility confined within the sphere rather than outside the sphere so that as history proceeds utility will always be outstripping disutility. So there may be some place for Vallentyne's intuitions to be fitted into an extended utilitarianism that handles infinite utilities.

Here we will look at a particular point that will have to be kept in mind if such an account is to be developed and one desires utilitarianism to give an absolute standard for right action. According to relativity theory, just as motion is relative to a reference frame, so is the simultaneity relation. The possibility is raised that we may encounter situations such as the following. Let F_1 and F_2 be different inertial reference frames. We will consider two infinite sequences of events C_1, C_2, C_3,. . ., and D_1, D_2, D_3, . . . It could happen that from both $F_1$and F_2, for $i < j$, C_i occurs before C_j, and D_i occurs before D_j. Yet it might also be the case that relative to F_1, for each i, C_i is simultaneous with D_i, and relative to F_2, D_1 precedes C_1, and, for each i, D_{i+1} is simultaneous with C_i. Now imagine that each of the Cs has a positive utility of one unit; D_1 has a negative utility of minus one half, and each of the other Ds has a negative utility of minus one. Suppose action a5 brings about both the C- and D-sequences, and action a6 would have prevented both the C- and D-sequences from occurring. And finally suppose that the only utilities involved in deciding whether a5 or a6 is better are those involved in the C- and D-sequences. In this case, applying PMU^* and U from F_1 tells us that a5 is better since after C_1 occurs there will always be a positive net utility of one half. On the other hand, from F_2 it will turn out that a6 is the better action since once D_1 occurs there will always be a negative net utility of one half. Thus the rightness or wrongness of an action will be relative to an inertial reference frame! This may strike one as counterintuitive. Perhaps if Vallentyne's principle is to be applicable it will be so only when the results are the same from every reference frame.[6]

University of Louisville

Received October 1993

[6] My thanks to Robert Cain for his comments.