Infinite Utility and Temporal Neutrality

Peter Vallentyne

Suppose that time is infinitely long towards the future, and that each feasible action produces a finite amount of utility at each time. Then, under appropriate conditions, each action produces an infinite amount of utility. Does this mean that utilitarianism lacks the resources to discriminate among such actions? Since each action produces the same infinite amount of utility, it seems that utilitarianism must judge all actions permissible, judge all actions impermissible, or remain completely silent. If the future is infinite, that is, the prospects for utilitarianism look bleak.

This very interesting criticism was made by Mark Nelson in his 1991 "Utilitarian Eschatology".[1] In response to this criticism of utilitarianism, I argued in my 1993 "Utilitarianism and Infinite Utility"[2] that utilitarianism had the resources for dealing with at least some cases of infinite utility. More specifically, I defended the following two principles as being part of the "spirit" of utilitarianism:

U: An action is permissible just in case no alternative action produces more utility.

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'.

The first principle is standard. The second principle was intended to be understood as a technical definition of "more utility" for the purposes of utilitarian theory. For the usual finite cases it agrees exactly with the usual approach of comparing totals. For infinite cases, however, it can distinguish among two actions each of which produces an infinite amount of utility. It says, for example, that an action that produces 2 units of utility at each time is better than an action that produces 1 unit of utility at each time -- even though both produce the same infinite amount of utility.

Of course, cases where one action produces more utility than a second at each time are going to be rather rare. PMU* is not limited to such cases, however; it also has bite in cases such as the following:

                        Time
a1:   1, 1, 1, 1, 2, 2, 2, 2, 2, ...
a2:   3, 3, 3, 3, 1, 1, 1, 1, 1, ...

For although the cumulative utility of a1 up to and including the fourth time is only 4 utiles and a2's is 12, a1 starts catching up after that, ties a2 in the 12th time, and stays ahead after that. Thus, PMU* judges a1 as producing more utility than a2.

I claimed that PMU* better captured the "spirit" of traditional utilitarianism that the usual sum total view -- on the grounds, for example, that almost anyone inclined to defend traditional utilitarianism would, upon reflection, want to hold that 2 utiles at every time is indeed better (greater, in the relevant technical sense) than 1 utile at each time.

In their 1994 paper "The Problem of Endless Joy: Is Infinite Utility Too Much for Utilitarianism?"[3] Jorge Garcia and Mark Nelson raise several objections to PMU*. Below I shall reply to these objections, but first I want to provide a bit more background to the general issue.

First, infinite futures do not automatically give rise to the above problem. Infinite futures generate infinite utility only under certain conditions. These conditions are laid out in each of the above articles, and I won't repeat them here.

Second, the problem can arise, not only for utilitarianism (with its welfaristic theory of individual goodness), but also, for any consequentialist theory that ranks states of affairs on the basis of the sum-ranking of individual goodness.

Third, although neither Mark Nelson, nor I, was initially aware of this, economists have been working on this problem since the 1960s.[4] Although work on the problem continues, the predominant view seems be that something like my PMU* is a correct principle. I mention this primarily to make readers aware of the interdisciplinary connection -- although the independent support for PMU* is of course welcomed.

Fourth, although Nelson and I focussed on the problem of infinite utility that can arise if the future is infinitely long, the problem is really much more general. The problem can arise if the past is infinitely long. It can also arise if time is finite but space, and the number of people, is infinite. And a comparable problem can arise in assessing uncertain states of affairs (where the outcome of an action depends on what the state of the world is) where there are an infinite number of possible states of the world. In "Infinity and Finitely Additive Value Theory" (unpublished) Shelly Kagan and I defend a much more general principle that covers all these cases. We also discuss the connection of our general principle (and PMU* by association) with non-standard math, which recognizes and allows arithmetic for non-standard infinite numbers.

Fifth, and finally, I want to make a retraction. In my original paper, I claimed that PMU* captured the spirit of traditional utilitarianism better than the sum-total view applied to infinite cases. In his beautiful "Infinite

Utility"[5] James Cain has shown that strictly speaking I was wrong. The problem is this: PMU* is formulated in terms of total utility <u>at a time</u>, whereas traditional utilitarianism takes <u>people</u> to be the basic bearers of utility. For finite cases, and indeed most infinite cases, this difference makes no difference to the utilitarian evaluation of actions. But in some special infinite cases it does. And where there is a difference, Cain rightly points out, the spirit of traditional utilitarianism goes with the person-centered approach rather than the time-centered approach of PMU*.

This doesn't mean that the core idea of PMU* must be completely rejected. For PMU* is, I would still argue, a plausible principle for a time-centered utilitarian (or sum-ranking consequentialist) a theory. And furthermore, there is a person-centered cousin of PMU* that is plausible for traditional utilitarianism. To fully develop and defend this last claim would, however, lead us too far away from the Garcia and Nelson paper, so I must refer interested readers to my "Infinite Utility: Person vs Time Centered Approaches".[6]

So much for a rather lengthy background statement. I turn now to the criticisms of PMU* made by Garcia and Nelson. For these purposes I shall assume for simplicity that treating times as the locus of value is unproblematic (since Garcia's and Nelson's criticisms do not depend on this issue).

One objection raised by Garcia and Nelson is that a defining feature of traditional utilitarianism is that it ranks the states of affairs produced by actions solely on the basis of the total utility they contain, whereas PMU* is sensitive to how utility is distributed over time. Consequently, the objection goes, PMU* can hardly be part of the spirit of traditional utilitarianism.

I acknowledge that for finite cases utilitarianism is deeply committed to ranking states of affairs on the basis of total utility. And for such cases PMU* requires that this to be so. Indeed, for finite cases, PMU* is equivalent to the more usual sum-ranking criterion. For infinite cases, however, things are much less clear. For the infinite cases in question the totals are infinite, and it's by no means obvious that the spirit of traditional utilitarianism requires all actions producing an infinite total to be ranked equally. For utilitarianism has been developed and defended almost exclusively with the finite case in mind. Consequently, it's not clear what the core commitments are in infinite cases. This is a matter, I claim, that must be explored rather than assumed.

A second and closely related objection is that utilitarianism is committed solely to maximizing utility, whereas PMU* is committed to producing utility as quickly as possible. As an analogy, Garcia and Nelson ask us to consider versions of these principles with books read replacing utility. According to PMU* reading two books everyday produces "more

books read" than reading one book everyday. And yet surely, Garcia and Nelson claim, for the infinite case, no more books are read. Both courses of action lead to the same infinite number of books being read. PMU*, it is claimed, is incompatible with utilitarianism's exclusive focus on <u>really</u> producing more utility.

Of course, I don't deny that the mathematical sums are the same in the 2 vs 1 scenario. And I agree that PMU* says that 2 everyday produces more than 1 everyday. The question is whether utilitarianism is concerned with standard mathematical sums in infinite cases. I claim that it isn't, and that there is a respectable sense in which 2 everyday does produce more than 1 everyday. After all, every single day more books are read under the 2 book approach than under the 1 book approach. And for any finite set of days -- no matter how large -- the total books read is greater under the 2 book approach. Consequently, there is a sense -- admittedly one that goes beyond the standard mathematical sum sense -- in which the total books read under the 2 book approach is greater.[7] Utilitarianism is not logically committed to using only standard mathematical sums for the infinite cases. It is free to adopt whatever standards of assessment it feels are appropriate. And I claim that PMU* is an appropriate standard for utilitarian purposes.

Another closely related objection raised by Garcia and Nelson is that PMU* fails to respect utilitarianism's commitment to a form of impartiality

(sometimes called "anonymity", "equity", or "neutrality") according to which switching (permuting) the utility values of any two bearers of value (e.g., times or people) has no affect on how an action is assessed. The idea is that every bearer of value is treated the same: no one gets special treatment. Utilitarianism's commitment to this form of impartiality in the finite case is obvious: switching utility values between two people or times has no affect on the total, and thus has no affect on how actions are assessed. For the finite case PMU* also respects this form of impartiality, but for the infinite case, I shall now illustrate, it violates a certain strong form of impartiality.

Consider the following case:

a1: 1, 1, 1, 0, 1, 0, ...
a2: 1, 0, 1, 0, 1, 0, ...

PMU* judges a1 as producing more utility than a2 -- even though a1 is just a permutation of a2. The distribution of a1 can be obtained by permutation from a2 in two steps. First, take a2 and switch leftward (i.e., by permutation) all the 1's that have a 0 to the left. This yields <1,1,0,1,0,1,0,1....>. Then do it again, and you've got a1.

So PMU* is incompatible with a certain strong form of impartiality in infinite cases. But, I claim, the incompatible form of impartiality is not a

core commitment of utilitarianism. We need to distinguish between a finite and an unlimited form of the impartiality condition. The finite form says that no <u>finite</u> number of switches makes any difference. PMU* is fully compatible with this principle. The unlimited form of impartiality says that no number (finite or infinite) of switches makes any difference. PMU* is indeed incompatible with this principle, as the above example shows.

Now, Garcia and Nelson hold that utilitarianism's commitment to this unlimited impartiality is clear and deep, and so they conclude that PMU* is not compatible with the spirit of traditional utilitarianism. I deny that the commitment is either clear or deep. I deny that it is clear because I think that before careful analysis nothing is clear in infinite cases. Utilitarianism was developed and defended with the finite case in mind. How things go in the infinite case needs to be explored.

I deny that utilitarianism has any deep commitment to unlimited impartiality because this impartiality is incompatible with a principle which I do think is a deep commitment of utilitarianism. For unlimited impartiality is incompatible with the following principle:

Monotonicity: If every ultimate bearer of utility (e.g., person or time) bears <u>at least as much</u> utility if action a1 is performed than as if action a2 is performed, and if at least some bearers of utility bear <u>more</u> utility under a1

than under a2, then a1 produces more utility than a2.

This monotonicity principle uses a Pareto criterion as sufficient condition for producing more utility.[8]

The example given above illustrates the incompatibility. For unlimited impartiality requires that the two actions in that case be ranked the same, where as Monotonicity requires that a1 be ranked as producing more utility than a2.

So the incompatibility is not merely between unlimited impartiality and PMU*, it is also between unlimited impartiality and Monotonicity. And it seems pretty clear to me (although I agree that there is room for reasonable disagreement) that Monotonicity is a deeper and more fundamental commitment of utilitarianism. Consequently, PMU*'s incompatibility with unlimited impartiality is not troublesome.

So far we have considered three related criticisms of PMU*: that it is incompatible with sum-ranking, that it is not committed solely to producing more utility, and that it is incompatible with impartiality. I turn now to a very different sort of criticism make by Garcia and Nelson, namely the criticism that PMU* in conjunction with U is effectively extremely undemanding in that it would judge very few actions impermissible.

How effectively demanding PMU* and U are, depends, of course, on

what the world is like. If the future is only finitely long, or if the utility impact of any action dwindles away to 0 ultimately, then PMU* and U are equivalent to standard utilitarianism. Where the utility impact of actions is infinite, the demandingness of PMU* and U depends on how the utility impact varies over time. If, for example, the difference in utility impact of actions disappears after some suitably long finite period of time (e.g., if actions can only affect what happens for the next billion billion years), then PMU* and U will agree with standard utilitarianism (since actions will be ranked on the basis of their impact during that period) and be just as demanding.

Of course, under a broad range of conditions PMU* will be effectively less demanding than traditional utilitarianism is in finite cases. For often there will be lots of feasible actions none of which produces more utility in PMU*'s sense than any other. It's not clear, however, how utilitarians should react to this possibility of decreased demandingness. After all, PMU* and U require agents to do the best they can in producing utility. And they clearly do judge some actions as wrong, namely those that produce less utility (e.g., 1 at each time instead of 2). Utilitarianism, I'm suggesting, is deeply committed to requiring that agents do the best they can, but it is not <u>deeply</u> committed to there being few permissible actions. Indeed, it's not even clear that utilitarianism is committed at all to there

11

being few permissible actions in a given choice situation. If a high percentage of the feasible alternatives are maximally good, then utilitarianism has no problem judging them all permissible.

So I deny that there is any deep commitment internal to utilitarianism that requires that most actions of a given choice situation be judged impermissible. Of course, from an external viewpoint -- one concerned simply with assessing how plausible a moral theory is -- the lower level of effective demandingness of PMU* and U may indeed cast some doubt on them. But once again, I claim, this is not clear. It all depends on what the appropriate level of demandingness is in such infinite cases. Although (because I am not a utilitarian!) I fully agree that PMU* and U do not impose that right sorts of demands on agents (e.g., to keep promises, to refrain from harming others, etc.), those problems are shared with traditional utilitarianism in the finite cases. But it's not clear that there is any problem that is not already present in the finite cases for the standard sum-ranking versions.

In any case, my most basic claim is that PMU* conjoined with U is more plausible than the standard sum-ranking approach.[9]

Notes

1.Nelson 1991

2.Vallentyne 1993

3.Garcia & Nelson 1994

4.See, for example, Von Weizsacker 1965; Diamond 1965; Gale 1967; Koopmans 1972; Campbell 1985; Epstein 1986; Lauwers 1993; and Van Liederkerke 1994a. Thanks to John Broome and Luc Van Liederkerke for bringing the existence of this literature to my attention.

5.Cain 1995

6.Vallentyne 1995b

7.The claim that more books are read by the two book per day policy is even easier to defend if one assumes that the books are numbered, and that each day the two book policy involves reading the next two unread books, whereas the one book policy involves reading the next <u>even-numbered</u> unread book. For in this case, the two book policy involves reading every book that the one book policy does, but not vice-versa.

8. Economists have recognized the incompatibility of infinite impartiality and some sort of monotonicity principle. See, for example, Van Liederkerke 1994b; and Lauwers 1993. I take the example and the statement of the incompatibility directly from the former.