

“Infinite Utility: Anonymity and Person-Centeredness”, *Australasian Journal of Philosophy* 73 (1995): 413-420.

Peter Vallentyne, Virginia Commonwealth University

1. INTRODUCTION

In 1991 Mark Nelson argued that if time is infinitely long towards the future, then under certain easily met conditions traditional utilitarianism is unable to discriminate among actions.¹ For under these conditions, each action produces the same infinite amount of utility, and thus it seems that utilitarianism must judge all actions permissible, judge all actions impermissible, or remain completely silent.

In response to this criticism of utilitarianism, I argued that utilitarianism had the resources for dealing with at least some cases of infinite utility.² More specifically, I defended the following principle as being part of the "spirit" of traditional (total) utilitarianism:

PMU*: An action a_1 has better consequences than an action a_2 if and only if there is a time t such that for any later time t' the cumulative amount of utility produced by a_1 up to t' is greater than that produced by action a_2 up to t' .³

For the usual finite cases, this principle agrees exactly with the standard 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 has better consequences than an action that produces 1 unit of utility at each time -- even though they 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 having better consequences than a2.

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

Two responses to my view have since been given in this journal. Luc Van Liedekerke has argued that PMU* violates an anonymity (neutrality, impartiality) condition that is central to utilitarianism.⁴ And James Cain has argued that PMU* is not person-centered, and that being so is central to utilitarianism. Here I will reply to the criticisms. Basically, I will acknowledge that each is right about an important point, but deny that this establishes that the core idea underlying PMU* is less plausible than the standard sum-ranking versions of utilitarianism.

2. VAN LIEDERKERKE AND TEMPORAL NEUTRALITY

Van Liederkerke formulates two conditions, which he claims are central to utilitarianism, and shows that in the infinite case these two conditions are sometimes incompatible.⁵ One condition is:

Monotonicity: If every ultimate bearer of utility (e.g., person or time) bears at least as much utility under action a1 as under action a2, and if at least some bearers of utility bear more utility under a1 than under a2, then a1 has better consequences than a2.

This, I agree, is a core commitment of utilitarianism.

The other condition is:

Strong Anonymity: If the utility profile (i.e., a specification of utility for each ultimate bearer of utility [e.g., time or person]) of one action is simply a permutation (i.e., rearrangement by interchanging numbers) of the utility profile of a second action, then the two actions are equally good.⁶

This condition also goes under the names of "impartiality", "equity", and "neutrality". The idea is that every bearer of value is treated the same: no one gets special treatment. All that matters are the utility numbers; it doesn't matter which bearer has which number. For example, 2 for i and 1 for j is equally as good as 1 for i and 2 for j. Utilitarianism's commitment to this form of anonymity 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.

Van Liederkerke shows that these two conditions are not always jointly satisfiable in the infinite case by means of the following beautiful example:

a1: 1, 1, 1, 0, 1, 0,

a2: 1, 0, 1, 0, 1, 0,

Assuming (as I do throughout) that the bearers of value (e.g. times or people) are the same (and aligned) for each action, Monotonicity directs that a1 be judged as having better consequences than a2 -- since a1 produces more utility than a2 for the second bearer (1 vs 0), and produces an equal amount for all other bearers.

Strong Anonymity, however, directs that the consequences of a1 and a2 be judged equally good. For 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 (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 Monotonicity and Strong Anonymity are not always jointly satisfiable. In his article Van Liederkerke does not take a stance on which condition utilitarians should reject. His point, an important one, is that for the infinite case utilitarians can't have both.

So which condition should utilitarians reject? I claim that Strong Anonymity is less central to the spirit of utilitarianism. Monotonicity, after all, just says (for the case of people) that if everyone is at least as well off, and some are better off, then the net result is better. Pretty hard for

utilitarians to argue with that!

Strong Anonymity, on the other hand, is not so central. To see this, it is useful to contrast it with the following weaker condition:

Finite Anonymity: If the utility profile of one action is simply a finite permutation of the utility profile of a second action, then the two actions are equally good.

This is exactly like Strong Anonymity except that it has bite only where the utility profile is a finite permutation of another -- i.e., only where one profile can be transformed into another by means of a finite number of binary permutations. Unlike Strong Anonymity, it is silent where only an infinite number of binary permutation can transform one utility profile into another. Thus, Finite Anonymity is silent in the above example - since the transformation requires an infinite number of binary permutations.

Monotonicity and Finite Anonymity, it is easy to show, are always jointly satisfiable.⁷ And PMU* satisfies the temporal version of each. If one action has better consequences for some times and as least as good for all times (the antecedent of Monotonicity), then PMU* judges the first action as having better consequences than the second (as required by Monotonicity) - since any interval including at least one of the better times will have a greater cumulative utility. PMU* also satisfies the temporal version of Weak Anonymity. For, if the utility profile of one action is simply a finite permutation of that of a second action (the antecedent of temporal Finite Anonymity), then the two utility profiles differ in utility values for at most a finite number of times. Consequently,

any set of times that includes all of the finite number of bearers where the utilities are different will have the same cumulative utility under both actions (since the utility profiles are permutations of each other). Thus, PMU* judges the consequences of the two actions as equally good (as required by Weak Anonymity). (PMU* violates Strong Anonymity - since in the above example PMU* judges a1 as having better consequences than a2, but Strong Anonymity directs that their consequences be judged equally good.)

Utilitarians cannot endorse both Monotonicity and Strong Anonymity. I claim that they should reject the latter, and endorse the weaker Finite Anonymity. Unlike Strong Anonymity, Finite Anonymity allows (and PMU* requires) that infinite permutations that have the effect of net shifts in utility to earlier points in time can make a difference in the assessment of a utility profile. It allows, that is, that where the total utility is infinite, the pattern of distribution of utility can serve a tie-breaker. Accepting Monotonicity and Finite Anonymity, and rejecting Strong Anonymity, is, I claim, the most plausible way of extending the spirit of utilitarianism to infinite cases.⁸

In sum: Van Liedekerke has very effectively shown that utilitarians cannot endorse both Monotonicity and Strong Anonymity when there are an infinite number of bearers of value. I have argued utilitarians should reject Strong Anonymity - although they should endorse Finite Anonymity. Given that PMU* satisfies the temporal version of both Monotonicity and Finite Anonymity, and that the standard sum-ranking approach violates Monotonicity, this provides further support for PMU*.⁹

3. CAIN AND PERSON-CENTEREDNESS

In response to my article, James Cain has provided a beautiful counterexample to my claim that PMU* captures the spirit of traditional utilitarianism better than the sum-total view applied to infinite cases.¹⁰ He is right, and I was wrong. But that doesn't mean that the core idea of PMU* should be rejected. I shall argue below that PMU* is plausible principle for a certain form of utilitarianism (admittedly not the usual or traditional form) and that a cousin of PMU* is plausible for traditional utilitarianism. First, however, I shall lay out what I take to be the core of Cain's criticism and the insight that it provides.

The problem Cain raises is this: PMU* is formulated in terms of total utility at a time, whereas traditional utilitarianism takes people 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. PMU* is not part of the spirit of traditional utilitarianism because it is time-centered. Let me explain in more detail.

Consider the following case, which is a modification of an example provided by Cain. Suppose that starting on a certain day and continuing forever exactly one person is born each day (and no one exists before that), and each person lives exactly three days. Furthermore, the n -th person born has 2^n units of happiness on the first day of her life, and -2^n units for each of the remaining two days of her life. Thus the situation looks like this:

<u>Person</u>	<u>Time</u>								<u>Total</u>	
1	2	-2	-2	-	-					-2
2	-	4	-4	-4	-	-				-4
3	-	-	8	-8	-8	--				-8
.	-	-	-	16	-16	-16	-	-		-16
.	-	-	-	-	32	-32	-32	-		-32
.
.
Total	2	2	2	4	8	16

In this example the total utility at each time is positive (since $2^n - 2^{n-1} - 2^{n-2} > 0$) and increasing. Consequently, PMU* would judge that an action with the above results produces has better consequences than an action that produced 0 utility for each person at each time. But this is incompatible with the spirit of traditional utilitarianism: for every single person has negative life-time utility (since $2^n - 2^{n-1} - 2^{n-2} < 0$). Consequently, traditional utilitarianism would not judge the above result as better than a "zero" result.

This example shows that for certain sorts of infinite cases it matters very much what the basic bearers of value are for principles like PMU*. For PMU* -- with its time-centered evaluation -- is plausible only if times are the basic bearers of value. Any principle that distinguishes between actions both of which produce the same infinite utility must be appropriately sensitive to what the

basic bearers of value are.

The problem with PMU* is that it privileges times over other possible bearers of value. There is, however, a more general basic idea underlying it, namely the following:

Core Idea: If the bearers of value that exist or that will exist are exactly the same if action a1 is performed as if action a2 is performed, and if there is a finite subset of these bearers of value such that for any superset the total utility produced for this superset by a1 is greater than that produced by a2, then a1 has better consequences than a2.¹¹

This principle is similar to PMU* but with sets of bearers of value as the locus of reference rather than time intervals. It replaces the reference to cumulative (total) utility for all longer encompassing intervals with a reference to total utility for all larger encompassing sets of bearers of value.¹² This core idea is also different in that it only states a sufficient condition for having better consequences (as opposed to PMU*'s necessary and sufficient condition). And finally, the core idea makes something explicit that was only implicit in PMU*: the principle only applies where the bearers of value (e.g. times or people) are exactly the same in the two situations. Where the bearers are times, this specification seemed unnecessary, since it seemed unlikely that the performance of an action can affect what times there are in a world. But in the case where people are the loci, the performance of an action can surely affect which bearers exist. And, as explained in the next footnote, in such cases the core idea would be implausible without the qualifying condition of there being the same people.¹³

This core idea is admittedly a very weak principle. It is silent whenever the existing bearers of value (e.g. people) are different for two actions. And even when the bearers are the same, it is very often silent. In some intuitive sense it will be silent more often than PMU*, since it will more often be the case that actions affect what people exist than that they affect what times exist. Consequently, it is less effective in answering Nelson's objection that utilitarianism loses its discriminatory powers in cases of infinite futures. Nonetheless, it provides at least a partial response. For it shows that utilitarianism does not lose all its discriminatory power in such cases.

To see this, consider the three examples considered above. The first is where one action produces 2 utiles for each bearer of value (e.g., time or person) and the other action produces only 1 utile. Assuming that the existing (present and future) bearers of value are the same for the two actions, the core idea judges that the first action has better consequences than the second (since a1 produces more utility for any finite set of bearers). And the core idea also judges that a1 has better consequences than a2 in the $\langle 1,1,1,1,2,2,2,2,\dots \rangle$ vs $\langle 3,3,3,3,1,1,1,1,\dots \rangle$ example (since any set of more than 12 bearers of value will have a greater total under a1).

Finally, in Cain's example, the core idea's judgements depend on what the basic bearers of value are. If people are the basic bearers, then the core idea judges the action in question as having worse consequences than an action producing no utility for any one (since every single person has negative utility under the action in question). If times are the basic bearers of value, then the core idea judges the action in question as having better consequences (since every single time has positive utility). And finally, if person-stages (people at a time) are the basic bearers of value, then the core idea is silent (since, no matter what finite set of bearers of value one considers, there are

supersets relative to which the total utility produced by the action is positive, and other supersets relative to which it is negative).

The core idea's sensitivity to what are the basic bearers of value is somewhat strange. In the finite case, if the utility values are fixed for each person at a time (for example), it doesn't matter whether the basic bearers of value are people or times. The total utility will be unaffected. But in infinite cases PMU* and the core idea are highly sensitive to what the basic bearers of value are. That, I take it, is one of the central insights that Cain's example provides.

One might be inclined to say "so much the worse for PMU* and the core idea", but that won't do -- at least not as long as one is committed to the principle that if every single bearer of value bears more utility under a_1 than under a_2 , then a_1 has better consequences. For in Cain's example this principle is sensitive to what the bearers of value are in exactly the same manner as the core idea is. And this principle, which is a weakened version of the core idea, is at the very heart of utilitarian thinking: if every single bearer of value bears more value under a given action than under another, then the given action is morally preferable.

In sum: Cain is right that time-centered PMU* is not compatible with the (person-centered) spirit of traditional utilitarianism. Nonetheless, the bearer-centered core idea articulated above, applied to people, is part of this spirit. The core idea admittedly has less discriminatory power than PMU*, and it thus less effective in responding to Nelson's original criticism. Still, the most plausible versions of utilitarianism -- even if subject to Nelson's and many other criticisms -- will incorporate the core idea to help deal with infinite cases.¹⁴

Notes

1. %Nelson 1991%

2. %Vallentyne 1993% Shortly before the present article went to press I discovered (via Luc Van Liedekerke's unpublished "Sacred") that the problem that arises when time is infinitely long had been previously discovered and addressed in %Seegerberg 1976%. The article by Seegerberg and I, and the article by Seegerberg and I, provide some very interesting equivalent formulations, and provide somewhat more questions about the difference between being equally good and for being at least as good. It's most unfortunate that this very interesting article has been overlooked.

3. In my original article I used the phrase "produces more utility" rather than "has better consequences", but I made it clear that it was to be understood as a technical concept for the purposes of utilitarian theory (as opposed to a standard mathematical concept). The phrase "has better consequences" is confusing.

4. %Van Liedekerke 1995%

5. See also %Lauwers 1993% for a statement of the incompatibility.

6. Van Liedekerke calls this condition "anonymity", but since I want to distinguish it below from a weaker version, I call it "strong anonymity".

7. Proof: If the antecedent of Finite Anonymity is satisfied, the utility profile of one action can be transformed into that of another action by a finite number of binary permutations. Thus, the two actions must produce the same utility for all except a finite number of bearers. And since the utility profiles are finite permutations of each other, for this finite set of bearers, each utility number that is produced under one action is also produced under the other action - although for a different bearer. And since this is a finite set, there is a maximal utility value, and it is produced by one action and produced by the other action for some other bearer. Since we are considering only bearers for which the utility values are different under the two actions, this means that at least one bearer (one with the maximal value) has more utility under the first action, and at least one

maximal value) has more utility under the second action. Consequently, Monotonicity is silent. So, Monotonicity and Weakly Pareto jointly satisfiable.

8. An alternative, of course, is to accept Strong Anonymity and a weaker version of monotonicity that is applicable only where there are a finite number of bearers of value. But this is a greater departure from the spirit of utilitarianism.

9. For more discussion of whether PMU* fails to be appropriately impartial (anonymous) see Garcia & Nelson 1994 and Vallentyne & Kagan 1994. The latter contains references to many articles by economists on the topic of infinitely long utility streams.

10. Cain 1995

11. The core idea is roughly the principle SBI1 that is defended in Vallentyne and Kagan "Infinite Value and Finitely Additive Value". The difference is that for simplicity the core idea leaves out a technical complication to deal with the possibility that, for all finite sets of bearers of value, one action produces a greater total than a second, but at the limit the actions produce the same finite total (as in $1+0+0+0\dots$ vs $1/2+1/4+1/8+\dots$). In speaking, the core idea is implausible for this reason, but it would take me too far astray to introduce the needed qualification. The core idea (SBI1 really) is a special case of a more powerful principle that Kagan and I defend in the above article. Our more powerful principle deals with a wide range of cases of infinite utility arising from infinite time, infinite space, infinite people, and even infinite states of affairs (uncertain states of affairs (where the outcome of an action depends on what the state of the world is)).

12. The replacement of the reference to time with a more general reference to sets of bearers of value requires that the notion of intervals (between all the bearers of value that lie between two specified bearers) not be presupposed. For although temporal bearers of value have intervals (between what), and thus can form intervals, many types of bearers of value, such as people, do not. People have no intrinsic order, so the core idea appeals to all possible finite sets of bearers of value rather than to intervals. Where the bearers of value are times, the

weaker than PMU*, since many sets of times are not intervals. This point, and its significance, is explained fully in Vallentyne and
and Finitely Additive Value Theory".

13. Suppose that people are the bearers of value, and that one action produces 2 units of happiness for each person, and the other
unit for each person. If the same people exist (present and future) no matter which of the two actions is performed, then the core
first action as better than the second (since for any finite set of people the total is greater under the first than under the second).
following variation: Suppose the second action involves creating 99 clones of each person, and, as before, produces 1 unit of hap
(including the clones). Surely traditional utilitarianism would not want to judge the first action as producing more utility than t
case, for each person with 2 units of happiness under the first action, there are 100 people with 1 unit of happiness under the seco
qualifier about the people being the same, the core idea would give this implausible result. With the qualifier, it is silent.

14. Thanks to James Cain, Brad Hooker and Mark Nelson for useful comments.