

JAMES L. HUDSON

THE DIMINISHING MARGINAL VALUE OF HAPPY PEOPLE

(Received 2 April, 1986)

ABSTRACT. Thomas Hurka has recently proposed a utilitarian theory which would effect a compromise between Average and Total utilitarianism, the better to deal with issues in population ethics. This Compromise theory would incorporate the principle that the value which an extra happy person contributes to a possible world is a decreasing function of the total population of that world: that happy people are of *diminishing marginal value*. In spite of its initial plausibility I argue against this principle. I show that the Compromise theory is actually no improvement over the two original versions of utilitarianism; in particular, it is subject to almost all the objections which are fatal to Average utilitarianism, and more besides. And I attempt to dispell the appearance that intuition supports the Compromise theory as against Total utilitarianism, by arguing that the latter's "Repugnant Conclusion," when properly understood, is not intuitively unacceptable. Total utilitarianism remains a plausible ethical theory, while both the Average and the Compromise theories should be definitely rejected.

In dealing with problems of population ethics, utilitarians have been torn between two versions of their theory – Average utilitarianism and Total utilitarianism – neither of which has seemed quite satisfactory. The differences between the ways in which these theories assign moral values to states of affairs can be represented as a difference in the importance placed respectively on quality and quantity of life. The Average theory places exclusive reliance on quality: one situation is better than another if the average happiness (that is, quality of life) of the people in the one is higher than that of the people in the other. The Total theory considers quantity in addition to quality: the average happiness must be multiplied by the number of people in order to get the moral value of the state of affairs, and thus a larger population is preferable to a smaller, *ceteris paribus*, provided the average is positive. According to Thomas Hurka, intuition suggests that the Average theory places too little importance on quantity, while the Total theory places too much.¹ Moreover, the deficiencies of the Average theory are especially glaring when it is applied to situations with very small populations, for then the moral importance

of extra happy people is relatively great; while the Total theory looks worst when applied to situations with very large populations, where it seems to exaggerate the value of excess population. The obvious solution is to find a middle ground between these two versions of utilitarianism, which will continue to take quality into account, while placing a variable importance on quantity. When population is small, extra happy people will be regarded in the manner of Total utilitarianism, as adding greatly to the moral value of the situation.² But when population is very large, something very near the Average point of view is appropriate, with extra happy people counting almost for nothing. Thus the Compromise theory will embody an important and little-noticed moral principle: The Diminishing Marginal Value of Happy People (DMVHP).

This principle is of a familiar sort, similar to the declining marginal demand for commodities in price theory, and the declining marginal utility of wealth in welfare economics. If it were correct, its discovery and explicit statement would mark an important advance in ethical theory. But I maintain that the principle is not correct: it has unacceptable implications, and receives from intuition only a spurious support. I propose to establish these contentions here. But first let us reconsider for a moment the strategy behind the construction of the Compromise theory.

We are to take two theories which by hypothesis are defective and find a middle ground between them which will be immune to the kinds of objections which tell against the original two. In following this strategy we must beware that our Compromise theory not instead fall prey to both kinds of objections. I shall contend that in fact the kind of compromise that would embody the DMVHP, while it would evade the principal objection to the Total theory, would be subject to the same criticisms that beset the Average theory; and that it would have an additional implausibility peculiar to itself.

In order to establish these claims I must canvass the most serious objections to the Average theory. But here a complication arises: as Hurka himself has shown, there is not just one Average theory – there are (by his count) eleven different ones which are worth distinguishing.³ Before a Compromise theory can be constructed, one of these Average theories must be selected to be one of the two poles

between which the compromise is strung.

So the structure of my paper is as follows. First I will sketch Hurka's taxonomy of Average utilitarianisms, and mention some of the principal objections to Average utilitarianism in general, and to some of the more interesting specific versions in particular. Second, I will take one of these Average theories (the most familiar one) and show that the theory that results from a compromise between this theory and Total utilitarianism will inherit the serious objections that confront the former. Finally, I will discuss the intuitive attraction of the Compromise theory and the DMVHP. I will try to show that bald appeals to intuition – on which Hurka rests his case – are inadmissible here; and that Total utilitarianism emerges from the discussion as the most promising of the theories discussed here.

I

The defining feature of an "Average" utilitarian theory (henceforth, "AvU") is that it imputes to "situations" or "states of affairs" a value which is the average of the levels of well-being (happiness, personal utility) of all people involved in the situation. Furthermore, since it is a form of consequentialism, it will always endorse that action which brings about a situation of the greatest possible value. We might express this by saying that actions are to be judged by their consequences. But there is room for further specification about just which situation is to be accounted the "consequence" of a given action. Is it to be (a) the whole possible world that would be actual if that action were done, or (b) the entire "future" of the world mentioned in (a) (i.e., the segment of the world in (a) consisting of everything later than the action), or (c) the temporal cross-section of the world in (a) taken just as the action is completed – the momentary world-state that would come into existence just as the action was done? Any of these things – whole worlds, "futures," or momentary world-states – can be evaluated by the average method. But which kind of evaluation is relevant to the moral judgment of *actions*?

The last of these, (c), is easy to dismiss. We must reject a theory which is so short-sighted as to evaluate an action by the situation which it brings about immediately, without regard for the further

future. But there *is* a way in which a consequentialist moral theory could make use of the evaluation of momentary world-states: it could use them to construct a different kind of evaluation of whole worlds, or of futures. Instead of taking the total all-time happiness of a world and dividing by the total all-time number of people to get the average, one could compute the averages for each moment in time, and then sum these to get the world total – or, if time is continuous, take the integral with respect to time of the momentary evaluations, considered as a function of time; and similarly for futures, with the integral taken over the limits “now” to “the death of the last person.”

To express this in mathematical notation, let U_t^i be person i 's utility level at moment t , and $(1, \dots, i, \dots, P)$ be all the people who ever live, each assigned his own natural number. The ordinary AvU theory prescribes the action for which

$$\frac{1}{P} \sum_{i=1}^P \int_{-\infty}^{+\infty} U_t^i dt$$

is greatest. (Note: dead or non-existent people are assigned $U = 0$.) With $t_0 = \text{now}$, the future-oriented version of this theory prescribes the action for which

$$\frac{1}{Q} \sum_{i=j+1}^{j+Q} \int_{t_0}^{+\infty} U_t^i dt$$

is greatest, where $(j+1, \dots, j+Q)$ are all the people who will live *from now on*. The new alternative form which uses average-type evaluations of momentary world-states would recommend the action for which

$$\int_a^b \frac{\sum_{i=1}^P U_t^i}{N_t} dt$$

was greatest, when $N_t =$ the number of people alive at t , and a and b are the first and last moments when at least one person existed, respectively. (I assume that at least one person exists at every moment

between a and b ; this simplifies the formula.) From this last theory we might derive a future-oriented version, requiring the maximization of

$$\int_{t_0}^b \frac{\sum_{i=j+1}^{j+Q} U_t^i}{N_t} dt;$$

but this would not yield any different judgments of actions, and so can be ignored.

We have thus discerned three distinct forms of AvU; these are three of Hurka's eleven, the ones he calls A1, A7, and A2, respectively. None of these three – indeed, none of the eleven – is free from very serious objections, yet one must be chosen to form the basis for our Compromise theory. I shall in the end adopt the most familiar of them, A1, for this purpose, and assume that my discussion of the resulting Compromise theory could be carried over to the compromises constructable out of the other ten AvU's. But I want first to give a general overview of the principal objections to all these forms of AvU, for I will contend that these objections will apply just as much to the corresponding compromise theories.

Let me begin by mentioning some serious theoretical weaknesses of any sort of AvU. First, it will require a sharp distinction between persons and lower forms of life, requiring the latter to be ignored. Lower animals simply cannot be counted in the average just as people are, and there does not seem to be any other way to include them into an averaging procedure.⁴ Thus torturing animals will contribute no intrinsic disutility to a situation. More important from the theoretical point of view, there seems to be a graduated progression from clear-cut persons to clear-cut non-persons, rather than any sharp distinction. Furthermore, some forms of AvU, in particular A1, require a sharp conception of personal identity. (All forms require at least that there be a definite number of people at any one time.) It will make a moral difference whether a given history is considered to be the history of a single person or of two different but closely related persons. But distinctions of this sort are also matters of degree, so any such theory is unsatisfactory (this is one of the themes of Derek Parfit's *Reasons and Persons*, although he says little about AvU).⁵

Finally, AvU does not allow us to evaluate situations in which no people exist. Some forms of the theory are embarrassed by the reflection that at some times there are no people, and all are embarrassed by the reflection that some possible worlds contain no people at all. A stipulation can be added to cover the case of zero people *ad hoc*, but this is theoretically unattractive.

A peculiar malady afflicts the future-oriented versions of the various Average utilitarian theories, such as A7. As Hurka notes ('More Average Utilitarianisms', Note 3) they give rise to inter-temporal inconsistencies: a certain action at a given time is evaluated differently depending on the temporal location of the evaluator. This is objectionable, because a moral theory should give us a single time-less evaluation of actions which is valid for everyone regardless of his location in time. Take for example, the addition to the population of a single person at time t , and suppose for simplicity that this person's net effect on the happiness of others will be zero, and that he will neither procreate nor affect others' procreation. Then whether his coming-into-existence is morally desirable depends on how he would affect the relevant future average utility. Perhaps at some earlier time t' it is true that his addition at t would raise future average utility – *future*, that is, with respect to t' ; then it would be judged good that he come into existence, bad that he not. But at some time t'' between t' and t a number of unhappy people may have died, so that from the perspective of t'' the addition of our hypothetical person at t would *lower* the future average utility (*future* with respect to t''). Thus his addition is judged desirable at t' and undesirable at t'' , and this not because of any new information (the deaths of the unhappy people may have been accurately foreseen at t'), but simply because of the change of temporal perspective. This cannot be correct.

Hurka offers a slightly different analysis of this type of case, in which his complaint is that future-oriented versions of AvU may require that at one time (t'') we try to prevent a state of affairs that at another time (t') we are required to try to bring about. This is not quite right, since all the forms of utilitarianism we are considering here are "objective" rather than "subjective" theories: they specify what one ought to do, not what one ought to *try* to do. Indeed, as a general rule one should (objectively) try to do an action only if he will

succeed – otherwise trying was probably a waste of effort, and so not utility-maximizing. In the specific example above, if I can do something at t' to *guarantee* that I will have a child at t , then I should do it. The issue will not arise at t'' , for even though at that time I recognize that it would be better that I *not* have a child at t , there is no longer anything I can do about it. On the other hand, if there is nothing I can do at t' to guarantee the result, and I foresee that at t'' I will be willing and able to prevent my having a child at t , then I had better not bother even trying to do anything at t' . But while I differ slightly in the analysis I agree with Hurka that the temporal inconsistency of future-oriented AvU is a very severe defect.

Another theoretical difficulty with AvU is the apparent arbitrariness of the choice among different versions. Dismissing for now the future-oriented versions, we find (as Hurka puts it) that A1 sums utilities over times and then averages over persons, while A2 averages over persons and then sums over times. Since this is simply a difference in the order of the operations, it is difficult to see a rationale for preferring one to the other. Yet they give different results in application; so a choice must be made – unless they are both rejected. Again, we could *sum over persons* and *average over times* – in either order. (Hurka, 'Average Utilitarianisms', Note 3). Or we could *average over both*, in either order. Thus Hurka obtains his multiplicity of versions of AvU. All are subject to serious objections; and an additional objection is that there is no particular reason to choose one over the others.

We have seen an objection specifically against the future-oriented versions of AvU; but most of the other forms (A2 is an exception) suffer from a different sort of defect. For according to these theories, whether or not we ought to perform a certain action depends on what human life was like in the distant past. Let us take A1 for example. The discoveries of archeologists may have a weighty bearing on our present-day decisions, altering completely our view of how we should act – not because they change our view of the present or the future, but because they alter our opinion about how numerous and how happy our distant human ancestors were; for this would alter our view of how our present actions might affect the all-time world-wide average happiness. This feature of A1 is very objectionable.

Furthermore we must, in principle, look beyond the familiar world

of persons – confined to the earth, or at most to the Solar System – and consider the possibility of intelligent moral agents existing elsewhere in the universe. According to all versions of AvU the levels of happiness attained by these beings, even if we never interact with them in any way, have a great effect on what we ought to do. If they are very happy compared with us, we ought to commit collective suicide (perhaps only by failing to procreate) so as to cease pulling down the world average utility. If they are much less happy than we are, we should be willing to produce children who will pull down the terrestrial average, for they may still increase the world average. All such consequences are very implausible, and apparently not to be remedied except *ad hoc*.

Finally let me mention a simplified version of Parfit’s “mere addition” problem. (See Note 4, above.) Suppose we have it in our power to create an absolute hermit – someone who never interacts with anyone else, not even indirectly – who is nevertheless relatively happy by ordinary standards. If we look at just these two alternatives – to create or not to create this hermit – which should we choose? Any form of AvU instructs us to make this decision by comparing the prospective utility-level of the hermit with some sort of average utility-level of other people. But the total isolation of the hermit largely nullifies the intuition that this average is relevant. The isolated hermit might as well be literally alone in the world, in which case it is obviously better to create him. AvU here demands a counter-intuitive procedure, and will often yield a counter-intuitive result.

A2 and the other forms of AvU which average across people at each moment of time suffer from an additional implausibility in regard to the hermit. Suppose we are faced with a world in which everyone has the same constant (positive) utility-level, but the size of the population varies in time. If the hermit’s utility-level would be higher than average, any form of AvU would instruct us to create him. But does it matter *when* he is created? His total isolation from the rest of the world suggests that this temporal factor is irrelevant; but according to A2 we should take care to create him when the population is smallest, so as to get the greatest (positive) effect on the average utility-level during his lifetime. Again, this result is not supported by intuition.

This catalogue of the defects of AvU makes no claim to completeness.⁶ My point in compiling it is just this: the Compromise theory that we construct using one of these versions of AvU together with TotU (Total utilitarianism) will, I claim, be found to be subject to all these same objections. In other words, the Compromise theory must be rejected for just the same reasons that a particular version of AvU must be. I will now illustrate this point, using A1 as my version of AvU.⁷

II

Let us construct a Compromise theory which will be like TotU when population is small and like A1 when population is large. An extra happy person added to a very small population will increase the value of the situation just about as TotU says; an extra happy person added to a very large population will have approximately the effect that A1 specifies. If we simplify by positing that all the relevant people – actual or merely possible – are at the same (positive) happiness level, then the extra person will increase the value of the small-population situation almost in proportion to his effect in increasing the population, while he will increase the value of the large-population situation almost not at all. In between the very small and the very large cases, the extra person's effect will be intermediate, steadily diminishing towards zero as total population increases. Hence the principle is that extra happy people have diminishing marginal value (DMVHP).

I will leave for later a discussion of the intuitive attractiveness of this Compromise theory. For now let me point out (1) that it does avoid the chief objection to TotU, but (2) that it falls prey to all the objections to A1.

The main complaint against TotU is that it leads to the Repugnant Conclusion,⁸ i.e., to the conclusion that a world containing only happy people in goodly numbers would be worse than a world with vastly more people at a much lower, barely positive, level of happiness. This conclusion follows only if TotU can be applied freely to situations with huge populations. But the Compromise theory resembles TotU only for small populations; for large populations it

yields quite different results, more like those of AvU. Now AvU originally recommended itself to philosophers precisely for its ability to avoid the Repugnant Conclusion and for related reasons. Clearly the Compromise theory is in good shape with regard to the Repugnant Conclusion.

But the listed objections to A1 do not similarly depend on applying both to large and to small populations. To be sure, Hurka claims that the counter-intuitive consequences of A1 are most glaring when population is small. But the more theoretical objections – that it is not forward-looking, that it depends on sharp notions of person and personal identity, etc., are not so dependent. The objections will apply also to the Compromise theory.

AvU, with its combination of theoretical weaknesses and (often) unintuitive results, is not a widely popular family of theories. (Admittedly, there are some distinguished advocates – for example, John C. Harsanyi.) The objections to it seem to me to be devastating. Then the Compromise theory cannot be much more attractive. At best, it removes a little of the conflict between A1 and intuition, by, in effect, applying A1 only to large-population cases. This helps to bring its results more into line with intuition, but it does nothing to remove the more theoretical objections.

And the Compromise theory is more complicated than either TotU or A1. Even if simplicity is not in itself a desideratum, this greater complexity produces an extra drawback. For a complete statement of the theory must tell us, as Hurka fails to do, what constitutes a “small” population and what constitutes a “large” one, and just how the value of an extra happy person falls as population increases. We know that in going from a population of zero to one, the happy person added counts at full value. As n increases without bound, in going from a population of n to $n + 1$ the extra happy person counts for an amount that approaches zero. But there is no hint in Hurka’s presentation of any way to establish the rate of decline, nor does any method independently recommend itself. Here is an apparent element of arbitrariness in the Compromise theory, which detracts from its acceptability.

In contrast TotU seems to be on much firmer ground. To be sure, there are weighty objections to any sort of utilitarianism; but we are

putting those aside in the present context. The only complaint specifically about TotU that arises in the context of population ethics is that it implies the Repugnant Conclusion; it produces just this one sort of unintuitive result. The prospect of producing a variant theory that avoids this result is indeed attractive; but not if the new theory has to bear the crushing weight of the objections to AvU, and more besides. The Compromise theory buys immunity from the Repugnant Conclusion at too high a price. Since the theory is not clearly an improvement on AvU, and seems clearly not to be an improvement on TotU overall, the original compromise strategy must be pronounced a failure.

III

In spite of all the foregoing, it must be admitted that the DMVHP exercises a powerful attraction. If nevertheless it cannot be made the basis of a successful Compromise theory, something must be said to reconcile us to abandoning it. Abandoning it while still finding AvU unacceptable, yet staying which the utilitarian tradition, means embracing TotU. But it is *very hard* to accept the Repugnant Conclusion. Let us look more closely into why this is so.

One reason may be that the whole setting of the population problem encourages us to take a sort of God's-eye view of alternative possible worlds, and to ask ourselves which world God would or should have made actual. In this context it is inevitable that we attribute a transcendent purpose to human life, which after all is being created to fulfill God's purpose. It is natural to conceive of this purpose in esthetic terms: God is to be thought of as a playwright, and human history is a drama of His creation. But however we think of the Divine purpose – whatever we think would be our purpose if we created the world – it is hard to see how mere numbers of people could contribute to it. Certainly if the purpose is esthetic, too large a cast of characters would spoil the drama, which requires rather a relatively small number of interesting characters interacting in interesting ways. To be sure, there is room for a “cast of thousands,” for thousands or millions or maybe even billions of “extras”; but the value of adding one more extra, when the plot has been decided on

and all the main parts filled, must be very minor. If an extra actor is added, with a guarantee that he will not get in the way of the others, then he can do no actual harm; but, after a certain point, he can do virtually no good, either.

But imagine now that the world drama as it actually has unfolded and will unfold, with its important events and interesting characters, were supplanted by vast hordes of peasant-like beings just well enough off not to be miserable, not doing anything interesting but merely struggling for a bare subsistence. This is one way of making the Repugnant Conclusion vivid, of filling in the details of a world which Total utilitarianism prefers even to a much improved version of our actual world. Now what sort of playwright would God be if he agreed with the Total utilitarians, and preferred to create the world teeming with peasants? Thus the Repugnant Conclusion cannot be correct, and extra people must be of declining marginal value.

If this sort of thinking lies behind the intuition, the latter should be dismissed. Theological ethics – the view that moral questions are on principle to be resolved by reference to God’s purposes – has been on the defensive ever since Plato’s *Euthyphro*: and besides, who would dare to assert that God’s purposes conform to human esthetic canons?

Besides irrelevant esthetics, irrelevant virtue-ethics may seem to support the intuition behind the DMVHP and against the Repugnant Conclusion. We may unconsciously reject the utilitarian conception that happiness is the good, feeling that something else – moral virtue, enlightenment, refinement and cultivation, or whatever – is more to be prized. Now our actual world – let alone an improved version thereof – probably contains much more of many of these other things than would a world teeming with peasants. Peasants tend to lack enlightenment and refinement; and scratching for a living in conditions of bare subsistence tends to promote a narrow focus on one’s own survival to the detriment of one’s moral virtue. So the teeming-peasants world may seem inferior to the actual one on this count.

But the point is not well-taken. First, the utilitarian can reply that these other things – virtue, etc. – are indeed of value, but that that value is instrumental rather than intrinsic. Virtue is to be prized simply because virtuous people tend to make mankind happier, while vicious people tend to make them unhappier. Then the value of

virtue can be accounted for even though *intrinsically* the only good is happiness. Second, the conception of the good is not really the point at issue with the DMVHP and the Repugnant Conclusion. True, the discussion is usually couched in utilitarian terms. But whatever one's conception of the good, he will have an analogous problem. There is another Repugnant Conclusion in which the actual world (assumed populated by virtuous people) is to be rejected in favor of one with vastly more people each of whom is just barely more virtuous than vicious; and there is another marginal principle – the Diminishing Marginal Value of Virtuous People (DMVVP). So this is not the place to quibble with the conception of the good which has been assumed for the sake of the argument.

Finally the DMVHP, as a way out of the Repugnant Conclusion, may seem attractive to the many people who are concerned in a more or less practical way with the overpopulation issue. Many fear that adding promiscuously to the population now may in the not-very-long run produce a catastrophe for mankind, of the sort predicted in the celebrated Club of Rome report.⁹ Whatever one's opinion of the accuracy of this prediction, it is easy to sympathize with the view that if the prediction is accurate it is vitally important that population growth be curtailed. And how does that sympathy comport with the Repugnant Conclusion, which seems to favor a population explosion?

Really, there is no conflict. The disaster predicted in the Club of Rome study is so enormous that a Total utilitarian would have excellent reason to wish to avoid it. After all, utilitarianism is against misery and short-sightedness on principle. What, indeed, is the population policy recommended by Total utilitarianism? Since the "total" in Total utilitarianism is *long-run*, the relevant question is: the production of how many babies in the near future will lead to the most total utility in the long run? The answer, for all I can see, might be either (a) more than, or (b) less than, or (c) the very same number as are going to be produced in the normal course of things. Any Total utilitarian persuaded by the Club of Rome's predictions, or by similar accounts of impending catastrophe, would presumably accept (b). Once it is seen that the Repugnant Conclusion concerns a long-run total, and is compatible with a concern to limit numbers of people in the near future, the DMVHP – put forward as a way out of the Repug-

nant Conclusion – should seem less attractive.

These reflections suggest that it is improper to rely on intuition in evaluating the Repugnant Conclusion or the DMVHP. Intuition is valuable insofar as it reflects common sense, which is always entitled to a certain respect. But there is no common sense about outré philosophical cases such as the choice between vastly different whole possible worlds. On such cases intuition can work only by unconsciously supplying a theoretical extension of common sense; and no great confidence should be placed in unconscious theorizing. Indeed, the most likely explanation for our having an intuition about a novel case is that we have misunderstood the case, conflating it with a more realistic sort of case; for the proper intuitive response to a truly bizarre case would be: “I’ll have to think about that!”

There may be other strains of thought tending to lend support to the DMVHP and to oppose the Repugnant Conclusion; but I know none that will stand examination.¹⁰ The fact that it yields the Repugnant Conclusion should not put us off from giving favorable consideration to TotU as a moral theory. On the other hand, possessing as it does most of the defects of AvU and few of the virtues of TotU, the Compromise utilitarian theory should be firmly resisted. While admiring Hurka’s representation of the charms of the DMVHP, let us not be seduced – its value is really submarginal.

NOTES

¹ Thomas Hurka, ‘Value and Population Size’, *Ethics* 93 (1983), 496–507.

² Hurka claims that the Compromise theory which *he* advocates gives *even more* importance to quantity in situations of very small population than does the Total theory. But unless there is some non-arbitrary unit for measuring personal utility, this claim is unjustified. Hurka could simply be interpreted as changing the units of measurement of utility, without making any substantive change.

³ T. M. Hurka, ‘Average Utilitarianisms’, *Analysis* 42 (1982), 65–69, and ‘More Average Utilitarianisms’, *ibid.*, 115–19.

⁴ For example, one could compute the average utility for each species, and add or otherwise combine all these figures into an overall measure of value (I owe this suggestion to Hurka). But this would be to attribute an overwhelming moral significance to the concept *species* – as opposed, say, to *genus*, or to *sub-specific variety* – which is quite implausible.

⁵ See Derek Parfit, *Reasons and Persons* (Oxford, 1984), sec. 143, ‘Why We Ought to Reject the Average Principle’. The basic reason offered is the “Mere Addition” paradox.

⁶ I haven’t mentioned the fact that some versions of AvU would have us kill off relatively unhappy people. See also J. McMahan, ‘Problems of Population Theory’, *Ethics* 92

(1982), 96–127, section VI; Bill Anglin, 'The Repugnant Conclusion', *Canadian Journal of Philosophy* 7 (1977), 745–54.

⁷ As Hurka has pointed out to me, in 'Value and Population Size' he actually presents a compromise between TotU and A2 (500f.); but the choice of A2 over, say, A1 seems arbitrary.

⁸ Parfit, *Reasons and Persons*, Chapter 17.

⁹ D. Meadows *et al.*, *The Limits to Growth* (New York, 1974).

¹⁰ I examine the repugnancy of the Repugnant Conclusion further in 'Ontological Ethics' (unpublished).

*Department of Philosophy,
Northern Illinois University,
DeKalb, IL 60115,
U.S.A.*