Broome on Moral Goodness and Population Ethics

Forthcoming in *Philosophy & Phenomenological Research*

Peter Vallentyne

1. Introduction and Overview

In an earlier book, *Weighing Goods*[1], John Broome gave a sophisticated defense of utilitarianism for the cases involving a fixed population. In the present book, *Weighing Lives*, he extends this defense to variable population cases, where different individuals exist depending on which choice is made. Broome defends a version of utilitarianism according to which there is a vague positive level of individual wellbeing such that adding a life with more than that level of wellbeing makes things morally better and adding a life with less than that level makes things morally worse. This version of utilitarianism avoids the extreme—but perhaps not all— forms of the repugnant conclusion that the usual total version faces.

As usual, Broome's work combines logical rigor with deep philosophical insight. There is much to learn from it. Nonetheless, I shall identify some problematic conditions used by Broome to derive utilitarianism and suggest that Broome's version of utilitarianism has implausible implications.

2. Prudential Goodness

Broome defends a utilitarian theory of moral (or general) goodness on the basis of various assumptions about moral goodness, prudential (or personal) goodness, and the relation between the two. In this section, I note my qualified agreement with his assumptions about prudential goodness. In the next section I shall question several of his assumptions about moral goodness.

I here grant Broome's claims that individual wellbeing (personal goodness, prudential goodness) is complete, transitive, continuous, strongly separable over states of nature, cardinally measurable, and interpersonally comparable in both units and zeros. Furthermore, zero wellbeing is, following Broome, set to be the value of a *constantly prudentially neutral life* (p. 254), which is a life at a constant level of wellbeing such that it is prudentially equally good to continue living at that level as to die.

Following Broome, I will also assume that the *wellbeing* of an individual in a world is defined only if the individual exists in that world. I reject, however, Broome's stipulation (p. 64) that a world can be *better for* an individual than a second world only if the individual *exists* in both worlds. There is, I agree, a narrow sense of "better for" (viz. "gives more wellbeing to)" in which this is true. There is, however, a broad sense of "better for" relative to which a world in which an individual does not exist can be better for her than a world in which she does exist. For example, a world in which Jones does not exist is better for her than a world in which she exists with extreme suffering and no offsetting benefits. If we were choosing between these two worlds solely on the basis of her prudential interests, we would choose the former world. In this broad sense, it is plausible—and I shall assume true—that a world in which an individual does not exist is equally good for her as a prudentially constantly neutral life. Given that the latter life has zero wellbeing, it follows that a world with positive wellbeing for an individual is better for her than a world in which she does not exist and the latter is better for her than a world in which she has negative wellbeing. In what follows, I shall use "better for" in this broad sense. The relevance of this will become apparent below.

3. Moral Goodness

Broome's argument for utilitarianism rests on the claim that moral (general) goodness is complete, transitive, continuous, strongly separable over states of nature, impartial, and satisfies the Principle of Personal Good (or Strong Pareto). The last two conditions are plausible in the fixed population case (although each, I would argue, requires some weakening in the variable population case). I would argue, however, that each of the other conditions should be rejected. I criticized these conditions in an article on Broome's earlier book and thus will comment on them only briefly here.[2]

Strong separability over states of nature (the sure thing principle) rules out risk aversion. Continuity rules out one consideration (e.g., Pareto superiority) being lexically prior to another (e.g., equality). Completeness rules out the possibility that some actions are incomparable (neither better, worse, or equally good) with others. Transitivity rightly rules out the possibility that z is at least as good as x when x is better than y and y is better than z. It also rules out, however, the possibility that x and z are incomparable. I see little reason to hold that moral goodness must satisfy these conditions. Indeed, I doubt very much that it does.

If these conditions are rejected, then Broome's argument for utilitarianism collapses. Nonetheless, in what follows, I shall, for the sake of argument, grant that, *in fixed population cases*, total utilitarianism is the correct theory of moral permissibility. I shall criticize the manner in which Broome extends this theory to the variable population case.[3]

4. Moral Permissibility and Moral Goodness

Broome defends a version of utilitarianism only as a theory of the moral *goodness* of actions. He leaves open how this relates to moral *permissibility*—except to insist that moral goodness is

sometimes positively relevant, and never negatively so, to moral permissibility. This makes it difficult for me to assess his theory, since I think that moral goodness of actions may have little content (be very incomplete). In what follows, I shall therefore assume what Broome (p. 31) calls *teleology*. This holds that an action is morally permissible if and only if it is sufficiently (e.g., maximally) morally good relative to the other feasible actions.

Broome quite explicitly refrains from endorsing (indeed seems to reject) teleology. Although I would reject teleology, it is only because I believe that the moral goodness of actions is too incomplete to ground the judgements of their permissibility. If we assume, as Broome does, that moral goodness is complete, then I see no reason to reject teleology. It leaves completely open how the moral goodness of actions is assessed. It need not be based on their consequences. Even if based on the consequences, it need not be solely based on the consequences for wellbeing. For example, I believe that rights considerations are relevant for permissibility. This is compatible with teleology as long as, as I believe, rights considerations are also suitably relevant for moral goodness of actions. The moral goodness of actions and the moral permissibility of actions are, I would argue, based on the same set of considerations. Moral goodness assesses how good an action is and moral permissibility assesses whether it is good enough, relative to its feasible alternatives, to be a morally acceptable choice.

In any case, I shall assume teleology when I apply Broome's version of utilitarianism. More specifically, I shall assume *maximizing teleology*, according to which an action is permissible if and only if it is a best feasible action. Broome may endorse maximizing teleology for the limited range of cases that I address (e.g., ones involving no killing, unfairness, or non-human goods). If so, then my objections apply to his theory of permissibility. If he does not endorse maximizing teleology for this limited range of cases, then the objections apply only to

those who combine his theory of moral goodness with teleology.

5. Critical-Level Utilitarianism

Broome defends a version of *critical-level (total) utilitarianism*, according to which one action is morally at least as good as another if and only if the total *standardized wellbeing* that it produces is at least as great as that of the other. An individual's standardized wellbeing is her wellbeing in excess of some specified (positive) *morally neutral level* (with shortfalls treated as negative numbers). For example, consider two actions that produce wellbeing distributions of <1,2,3> and <*,2,3> respectively, where "*" designates non-existence. Suppose that the morally neutral level of wellbeing is 2. The standardized totals of the two actions are thus 0 (-1+0+1) and 1 (0+1) respectively. Thus, the second action is judged morally better than the first—even though the first has a greater (non-standardized) total (6 vs. 5).[4]

The crucial idea of critical-level utilitarianism is that not all lives with positive wellbeing make things morally better. In particular, lives with positive wellbeing below the morally neutral level make things morally *worse* than their non-occurrence. This feature of critical-level utilitarianism enables it to avoid extreme forms of the repugnant conclusion that standard total utilitarianism faces (since, if the morally neutral level is set high enough, adding the low level lives makes things morally worse).

Critical-level utilitarianism can, however, face the negative repugnant conclusion— according to which it is better to have a large number of people with lives full of suffering than to have a sufficiently larger number of people with positive wellbeing below the morally neutral level. This is because the people with positive wellbeing have negative standardized wellbeing and, if there are enough of them, the total standardized wellbeing will be lower than that in the

world in which all people suffer. For example, if the morally neutral level is 2, then the total standardized wellbeing $\langle -1, -1 \rangle$ is -6 and that of $\langle 1,1,1,1,1,1,1 \rangle$ is -7.

Broome argues that the morally neutral level is *vague* and that, if the morally neutral level is set at the right level (neither too high nor too low) and is suitably vague, then the resulting theory can to a large extent avoid both problematic conclusions. I shall suggest, however, that, assuming maximizing teleology, it still has implausible implications for permissibility. To a very large extent, I will here be drawing the insightful work of Melinda Roberts.[5]

Critical-level utilitarianism entails that adding a person to the world with positive wellbeing but below the morally neutral level is always at least vaguely wrong (i.e., wrong on at least some admissible ways of sharpening the specification of the neutral level). Such a person has at least vaguely negative standardized wellbeing and adding her to the world thus at least vaguely reduces the total standardized wellbeing. I agree that it is sometimes wrong to add a person with such a life to the world—for example, when one could add that very same person with an even better life with no negative effect on anyone. I deny, however, that it is evenly vaguely wrong to add a person to the world *with a positive and best feasible life for her*—when no one else could be made better off. In such a case, no one is wronged and hence, I claim, the action is permissible. It is not even vaguely wrong to add such a life.

We can clarify the content of this objection by identifying the following conditions[6]:

P* (Permissibility): An action is permissible (in a given choice situation) if and only if it would (if performed) wrong no one.

This is the person-affecting restriction. It rules out impersonal wrongs (actions that are wrong without wronging anyone). This is a controversial assumption, but I believe that it is correct, and shall appeal to it below.

M* (Maximal Well-Being): A person is not wronged by an action if it is at least as good for her as any feasible alternative.

This rules out harmless wronging. Strictly speaking, I reject this condition because I believe that individuals have choice-protecting rights (e.g., physically forcing me to get a flu shot wrongs me even if on balance it is best for me). Here, however, I set this concern aside and work within a welfarist framework (as Broome does).[7]

To illustrate, suppose that there are only two possible people and only two feasible options: <*,5> and <1,5>. M* entails that no one is wronged by <1,5>—since each person existing therein is made as well off as possible. P* then entails that <1,5> is permissible. Critical-level utilitarianism, however, says that <1,5> is at least vaguely wrong if 2 is an admissible specification of the morally neutral level. That, I claim, is false. <1,5> is not even vaguely wrong.

A second problem with critical-level utilitarianism is that it can judge it wrong to *fail to create* a person with wellbeing determinately above the critical level—even where there is no effect on others. Suppose that the only two feasible options are <*,6> and <3,5>, where the morally neutral level is determinately below 2. In this case, the standardized total wellbeing of <*,6> is determinately below that of <3,5> (e.g., if the neutral level is 1, then the standardized totals are 5 vs. 6 and thus only the latter is judged permissible). I see, however, nothing wrong

with <*,6>. The only person who exists is the second person and no alternative is better for her. M* thus ensures that she is not wronged. The first person does not exist and thus, I claim, is not wronged either by <*,6>.

More generally, I claim that the following condition is plausible:

N* (Non-Existence): An individual is not wronged by an action if she would never exist (in past, present, or future) if the action were performed.

N* ensures that you never wrong a person for failing to bring her into existence. Obviously, many people reject this claim, but it seems very plausible to me. Critical-level utilitarianism, like standard total utilitarianism, violates this condition and that, I claim, is a problem.

A third problem with critical-level utilitarianism is that it can wrongly judge an action permissible even when there is some alternative feasible action that is at least as good for all those who would exist under both actions *and is better for some*. Suppose, for example, that there are only two feasible actions, <5,5,5> and <6,*,6>, where the morally neutral level is determinately below 2. Their respective standardized totals are thus at least 9 and 8 and thus the first is judged permissible. This seems quite mistaken. The second action is better for the first and third individuals (who exist under both actions) and the second person does not exist under the second action. It seems that priority should be given to making the first and third persons better off over *adding* the second person to the world.

When can generalize this point with the following condition:

D* (Deprived Gratuitously): An individual is wronged by an option X if (1) she exists in X, and (2) there is an option, Y, such that (a) everyone who exists in both X and Y is at least as well off in Y as in X, (b) Y makes her better off, and (c) everyone who exists in Y but not in X is as well off as is feasible.

In the above example, the first and third persons are wronged by <5,5,5>, since <6,*,6> makes each better off without disadvantaging anyone else who would exist. D* is a kind of optimizing condition that gives a certain priority to making people "who exist anyway" better off over adding people to the world.

It's worth noting that the above conditions are fully consistent with total utilitarianism for the fixed population case. They only do work in the variable population case. The following is an example of a theory that satisfies the conditions and is equivalent to total (and critical-level) utilitarianism in the fixed population case. This theory holds that an action is permissible if and only if (1) it is judged permissible by the above conditions or (2) it is not judged impermissible by the above conditions and it maximizes total wellbeing relative to those actions that are not so judged.[8]

The above conditions ensure that the standard form of the repugnant conclusion is avoided. For example, suppose that there are just two feasible options: <4,*,*,*,*> and <1,1,1,1,1>. The above conditions (like critical-level utilitarianism with a neutral level of 2) entail that only the former is permissible. Given P* (permissible iff wrongs no one), D* entails that <1,1,1,1,1> is impermissible (since the first person can be made better of without it being worse for any existing person) and M* and N* entail that <4,*,*,*,*> is permissible (since this is the best action for the first person and no one else exists).

The above conditions do, however, have a "repugnant-like" implication that many will reject. Suppose that there are just two (six person) feasible options: <4,*,*,*,*,*> and <*,1,1,1,1,1>. Condition D* is now silent (since there is no one who exists under both actions). Given P*, M* and D* entail that *both* options are permissible (since each is best for those that exist under it). I'm inclined to think that this judgement is correct. In particular, I see nothing wrong with creating people with lives barely worth living if the only alternatives are ones in which *different* people exist. In such case, no person would have a complaint. Broome disagrees (along with many others!). He holds that it is morally worse to add a person below the morally neutral level. So he holds that something can be morally worse even though no one has a relevant complaint. That's a kind of impersonal value that I would reject.

6. Conclusion

Broome offers a technically sophisticated and philosophically insightful defense of critical-level utilitarianism. I have flagged my rejection of most of the assumptions that he makes about moral goodness. I have further suggested that, for the variable population case, it is more plausible to take a person-affecting approach and ground moral permissibility directly in individual prudential goodness. Although I believe that these points can be developed into powerful objections, all I have done here is to highlight some of the big picture issues worth further exploration.[9]

---

[1] John Broome, *Weighing Goods* (Cambridge, MA: Basil Blackwell, 1991).

[2] See, Peter Vallentyne, "The Connection between Prudential Goodness and Moral Permissibility", *Journal of Social Philosophy* 24 (1993): 105-28.

[3] Broome extends utilitarianism to variable population cases in part by defending the separability of persons. The objections raised below cast doubt on this principle, but space limitations prevent me from developing this issue here.

[4] Critical-level utilitarianism is systematically explored by Charles Blackorby, Walter Bossert and David Donaldson, *Population Issues in Social-Choice Theory, Welfare Economics and Ethics* (New York: Cambridge University Press, 2005). One difference between the two views is that Broome argues that the morally neutral level is vague.

[5] Melinda Roberts, *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law* (New York: Rowman & Littlefield Publishers, 1998). Melinda Roberts, "A New Way of Doing the Best That We Can: Person-based Consequentialism and the Equality Problem", *Ethics*, 112 (2002): 315-350. See also Peter Vallentyne, "Critical Notice of *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law by* Melinda Roberts", *Nous* 34 (2000): 634-47.

[6] All the following conditions are taken from Roberts, *Child versus Childmaker*.

[7] M* is especially controversial in light of the non-identity problem. This is the problem generated by the fact that identity seems to be highly sensitive to origins (e.g., the identity of the particular egg and sperm that give rise to the individual). Thus, for example, if procreation now is the only way that a genetically severely disabled person, with a life worth living, will come into existence, then M* entails that she is not wronged by being brought into existence—even if waiting a day would have produced a perfectly healthy (but different!) individual. In *Child versus Childmaker*, Roberts gives a strong defense of the view that the non-identity problem is compatible with procreation wronging the procreated individuals even if they have lives worth

living.

[8] For more on developing theories of permissibility that satisfy the above conditions (with special emphasis on Paretian egalitarianism), see Bertil Tungodden and Peter Vallentyne, "Paretian Egalitarianism with Variable Population Size", in *Intergenerational Equity and Sustainability*, edited by John Roemer and Kotaro Suzumura, (Palgrave Publishers Ltd., 2006), ch. 11.

[9] For helpful comments, I thank John Broome, Brian Kierland, Melinda Roberts, Bertil Tungodden, and Paul Weirich.