

John Broome, *Weighing Lives*

Oxford: Oxford University Press, 2004. ix + 278pp.

Weighing Lives is a masterwork that everyone interested in ethics should read. In it, Broome develops a general theory for dealing with population problems, or problems in which how we act may affect who shall exist. On the way, he provides fascinating discussions of an array of issues from the general structure of value to the badness of death. The book involves some technicality, but everything is clearly explained and should be accessible to a wide audience. The writing is crisp and elegant; the progression of ideas is thorough and systematic; and many of the arguments are both very compelling and very surprising in their implications. While Broome's arguments are forceful, he is generally very aware of the problems they face, and in many ways he is his own best critic. In sum, *Weighing Lives* is impossible to read without wishing there were more books of its kind.

Broome's general problem is to compare the value of outcomes that differ concerning what people exist and how well they fare at any given time. More precisely, he approaches the problem of weighing lives as one of evaluating *two dimensional distributions*, where the two dimensions in question are those of people and of time. A two dimensional distribution is a specification of the *condition* of each person at each time, where the condition of a person at a time consists either in her level of well-being at this time (if she exists then), or in non-existence (if she doesn't). Broome's aim is to determine when one such distribution is better than another.

In his previous book, *Weighing Goods*, Broome tackled a special case of this problem. He considered cases where the distributions being compared differ not in terms of who is alive and for how long, but only in terms of how well people fare at various times. Concerning such cases, he presents an *intertemporal addition theorem* and an *interpersonal addition theorem*. The former theorem provides a formula for determining a person's level of *lifetime wellbeing* (or how well her life goes as a whole) in terms of her levels of *temporal wellbeing* over the course of her life (or how well she fares at various times). This theorem states that if two possible lives someone might lead have

the same duration, the first involves a greater level of lifetime wellbeing just in case it involves a greater sum of temporal wellbeing. Analogously, the interpersonal addition theorem states that if two distributions involve exactly the same people, the first is better than the second just in case it involves a greater sum of lifetime wellbeing. While Broome fully endorses the interpersonal addition theorem, he has strong reservations about the intertemporal addition theorem, and in *Weighing Lives* the latter serves only as a default assumption.

In *Weighing Lives*, Broome extends this account so as to compare the values of distributions that can differ in terms of who lives and for how long. To do so, he needs an account of the value of adding a life to a population, as well as an account of the value of extending a life. Providing these accounts proves no easy task. Much of the difficulty lies in defining a *neutral level for existence*, or a level of lifetime wellbeing such that adding a life at this level doesn't affect the value of a distribution. Similar difficulty lies in defining a *neutral level for continuing to live*, or a level of temporal wellbeing such that living an extra period of time at this level doesn't affect one's lifetime wellbeing. Broome argues that both these levels are unique. That is, there is only one neutral level for existence, labeled v , lives above which increase the value of a distribution, and lives below which decrease this value. Similarly, he argues that there is only one neutral level for continuing to exist, labeled μ , continuing to live at a level above which increases the value of one's life, and continuing to live at a level below which decreases this value. However, Broome argues that these two neutral levels are probably vague, and thus we cannot specify them with precision but can state only that each one lies within a certain range. Further, Broome argues that the neutral level for existence is independent of context, and is thus the same for every distribution. He also accepts the default assumption that the neutral level for continuing to live is independent of context, and is thus the same for all lives.

Broome then defines a person's *standardized lifetime wellbeing* as the degree to which her level of lifetime wellbeing exceeds v (the neutral level for existence). And he defines a person's *standardized temporal wellbeing* as the degree to which her temporal wellbeing exceeds μ (the neutral level for continuing to live). He argues for the

standardized total principle for distributions, a generalization of the interpersonal addition theorem that applies to distributions of varying populations. This principle states that one distribution is better than another just in case it involves a greater sum of standardized lifetime wellbeing. He likewise argues for the *standardized total principle for lives*, a generalization of the intertemporal addition theorem that applies to lives of varying length. It states that one life involves more lifetime wellbeing than another just in case it involves a greater sum of standardized temporal wellbeing.

The implications of the standardized total principle depend on the value of v . On the simple view held by many utilitarians, a life is neutral with respect to the total value of a distribution just in case, on average, it is lived at a level of temporal wellbeing at which it is neither beneficial or harmful to continue to live. That is, v is the level of lifetime wellbeing of a life whose average level of temporal wellbeing is μ . But on this simple view, Broome's principles imply Parfit's Repugnant Conclusion.¹ For Broome's standardized total principle implies that for any possible distribution, there is a better distribution involving a much larger population in which everyone's level of lifetime wellbeing is barely above v . Hence, on the simple view, for any possible distribution in which everyone leads a wonderful life, there is a better distribution in which everyone constantly lives at a level of temporal wellbeing just above μ , that is, just above the level where it ceases to be worthwhile continuing to live. One might try to avoid this conclusion by setting v at the level of a reasonably good life. But then Broome's principles imply what he calls the Negative Repugnant Conclusion: for any distribution in which everyone leads a horrendous life, there is a worse distribution involving a much larger population in which everyone's life is reasonably good.

Broome suggests that we can avoid both these conclusions by recourse to the vagueness of v . If the range of values over which v is vague includes reasonably high levels of wellbeing, we can avoid affirming the standard Repugnant Conclusion. And if this range also includes very low levels of wellbeing, we can avoid affirming the Negative Repugnant Conclusion. But this strikes me as an unsatisfactory response. For

¹ Derek Parfit, *Reason and Persons* (Oxford: Clarendon Press, 1984), 381-390.

we don't merely want to avoid *affirming* the Repugnant Conclusion: we want to *deny* it. And we can do so only if we can deny that v is very low. Likewise, we want to *deny* the Negative Repugnant Conclusion, and we can do this only if we can deny that v is reasonably high. But if the value of v is vague in the way Broome proposes, then we can make neither of these denials, and so we cannot deny either of these repugnant conclusions.

But the problem is still worse. Broome cannot even legitimately avoid *affirming* the Repugnant Conclusion, regardless of what value he assigns to v , and regardless of how wide we make the range over which the value v of is vague. For his standardized total principle for lives implies that for any given life, however wonderful, a sufficiently long life will count as better even if it is constantly lived at a level just above μ . Therefore, no matter how high we set the value of v , a long enough life lived at a level just above μ will have a lifetime wellbeing greater than v . Hence it follows from the standardized total principle for distributions that, for any possible distribution, there will be a better distribution in which everyone constantly lives at a level of temporal wellbeing just above μ .

Thus, given Broome's two standardized total principles, he cannot avoid affirming the repugnant conclusion. We might nonetheless reasonably accept Broome's theory if it rested on sufficiently secure foundations. However, in supporting his theory, Broome appeals to a number of questionable premises. First, he relies heavily on controversial principles he argued for *Weighing Goods* (especially the two addition theorems mentioned above, and his principles of personal and temporal good).² Hence any worries one might have about these arguments in *Weighing Goods* will carry over to *Weighing Lives*. And the latter work adds further contentious premises. For example, Broome argues that the neutral level for existence must be independent of the average wellbeing of the world's population, for otherwise it would depend on "how well off people were in the Stone Age," and such dependence, he says, is "incredible" (p. 194). But it's unclear whether this intuition can bear the weight that Broome's argument requires. For it's easy

² See Larry Temkin, "Weighing Goods: Some Questions and Comments," *Philosophy and Public Affairs* 23 (1994): 350-380.

to devise thought experiments in which many people have the contrary intuition that the neutral level for existence does depend on such remote conditions. And even if such dependence is hard to believe, it isn't nearly as incredible as the repugnant conclusion. And so here, Broome's *modus ponens* may be another philosopher's *modus tolens*.

In fairness to Broome, however, his aim is not to present a "comprehensive theory of weighing lives," but rather a "default theory" that can serve as a starting point for further inquiry (p. 128). And in relation to this aim, he succeeds admirably. I doubt anyone could have done a better job.

Jacob Ross

University of Southern California