

Summary

Philosophical Books, 48 (2007), pp. 289–9

We are very often faced with decisions that involve weighing the lives of some people against the lives of others, or weighing people's lives against other goods. For example, we often decide what risks to take with our own lives, and governments decide on priorities within the health service. Moreover, some of our decisions affect the existence of future people: a couple decides whether to have a baby, or a government sets its tax policy to encourage or discourage procreation. *Weighing Lives* develops a theory of the value – goodness – of people's lives. It asks how good it is for a person to continue living, and how good a person's existence is. The value of the lives involved is one of the central considerations that should determine life-and-death decisions. It is not the only consideration; fairness is another. But nevertheless the value of the lives is central, so my theory of value aims to contribute to the practical decisions.

My theory is more specifically a theory about the aggregation of value. I argue that the goodness of a person's life overall (which I call her 'lifetime wellbeing') is made up of the person's levels of wellbeing at the various times in her life (I call these her 'temporal wellbeings'), and my theory includes an account of how. It also includes an account of how different people's lifetime wellbeings together determine the overall value of the world. As I put it, this is a theory about the aggregation of value across two dimensions: across times within each person's life and across people.

My earlier book *Weighing Goods* also offered a theory about aggregation across these two dimensions, but it set aside the special problems that arise when the lengths of people's lives may vary, and when the number of people who make up the world's population may also vary. *Weighing Lives* takes up those problems. It adopts the theory of *Weighing Goods* and extends it. Aggregation is a major topic in economic theory, and some of my methods are adapted from economics.

Weighing Lives first (chapters 1 to 4) considers various issues concerning goodness and its position in an ethical theory. It explains the idea of a betterness ordering, examines the relation between goodness and rightness, compares the idea of goodness for a person with goodness overall, argues that betterness is necessarily a transitive relation, considers whether goodness can be relative to a point of view, and so on. It then goes on (chapters 5 and 6) to a discussion of the measurement of wellbeing. This is not a discussion of its empirical measurement, but a conceptual discussion of the quantitative concept of wellbeing. For a reason explained in my replies to the comments below, this question of measurement is an essential preliminary to the work of the book.

The topic of the book is the aggregation of value across two dimensions. It can be approached by two routes. We could start by evaluating the world as it is at each particular time – aggregating across people at each time – and then go on to aggregate across times to reach an overall evaluation of the world. Alternatively, we could start by evaluating each person's life – aggregating across times within the life – and then go on to aggregate across people to reach the overall evaluation of the world. But chapter 7 argues that the first of these two routes to aggregation will not work. It requires an assumption known as 'separability of times', which is not reliable. The book therefore takes the second route.

This means it has two tasks. It has to consider how the value of the world depends on the lifetime wellbeing of each person, and it has to consider how each person's lifetime wellbeing depends on that person's temporal wellbeings at all the times in her life. The book tackles these two tasks in turn.

For the first, I take my core theory of aggregation across people from *Weighing Goods*. It is a broadly utilitarian theory. Chapters 8 and 9 explain it, and place it within the new context of *Weighing Lives*.

Chapters 10, 11 and 12 consider a difficult problem that arises in the new context. Extended to this context, my theory implies that the value of adding a person to the population of the world depends on her level of lifetime wellbeing. Moreover, there is just one level of lifetime wellbeing (I call it the ‘neutral level’) such that adding a person at that level is equally as good as not adding her. If her wellbeing is above this level, adding her is better than not; if below, worse. But many people’s intuition is that adding a person to the population *generally* has neutral value: there is not just one neutral level but many. This is one consequence of what is often called ‘the person-affecting principle’. These three chapters explore various possible ways of accommodating this intuition within a coherent theory of goodness. I conclude in the end that the intuition is not consistent with any coherent theory. There is indeed just one neutral level. However, I think this level is vague.

Rejecting the intuition puts me in a position to develop my account of how lifetime wellbeing is aggregated across people. That happens in chapters 13 and 14. My method is broadly axiomatic: I derive the conclusion from a few plausible assumptions or axioms. Given the axioms, the conclusion is inevitable. The goodness of the world is the total of the lifetime wellbeings of the people, provided the scale of wellbeing is normalized to set the neutral level at zero.

Chapters 15, 16 and 17 undertake the book’s other task: the aggregation of temporal wellbeing across times within a person’s life. Here there is more choice; I do not have plausible axioms that determine a single conclusion. Instead, I adopt a ‘default theory’. It is the simplest plausible theory available. I would be happy to give it up if I found a good reason to do so, but I have not found a good reason.

Within the default theory, there is once again a sort of neutral level. This time it is a level of temporal wellbeing that is neutral for continuing to live. Consider the question of whether it is better for a person to continue living or to die. The neutral level of temporal wellbeing is the level such that continuing to live and dying are equally good for the person. Let us normalize the scale of temporal wellbeing to make this level zero. Then the default theory is that the goodness of a person’s life is simply the total of all her temporal wellbeings through her life, normalized in this way.

Finally, chapter 18 pulls the two dimensions of aggregation together to arrive at a comprehensive theory of value that encompasses both. It also discusses how the theory might be applied in practice.