

Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis

By MATTHEW D. ADLER

Oxford University Press, 2012. xx + 636 pp. £55.00

1. Introduction

Matthew Adler, a law professor at the Duke University, has written an amazing book in defense of a theory of wellbeing and a version of moral prioritarianism. He draws heavily on the framework and results of the theory of social choice, but he does so in an intuitive way with only a bare minimum of symbols. He also draws heavily on philosophical discussions of key issues, such as metaethics, personal identity, and wellbeing. The net result is technically rigorous and philosophically insightful.

Adler defends a very interesting extended-preference approach to the theory of wellbeing—one in which wellbeing is ratio-scale measurable and non-existence has value zero. Because he wisely allows that people may have different fully informed and rational extended preferences, wellbeing rankings need not be complete. Although this is a rich and important account of wellbeing, I shall focus, in what follows, on Adler's discussion of moral issues.

Adler defends prioritarianism as a moral *decision-making procedure*. Much of his defense, however, addresses the *criteria* of moral permissibility and goodness, and, since that is where my interest lies, I shall focus on his theory so understood. Thus, I shall not address his discussion of implementation issues, such as the estimation of wellbeing and how cognitive and motivational limitations affect the effectiveness of decision-procedures.

Adler assumes a maximizing welfarist consequentialism theory of morality and he limits his task to defending the most plausible version thereof. Moreover, he limits his attention to the

evaluation of *large-scale state policies*, and thus his view is not subject to the objection that maximizing consequentialism demands too much sacrifice from individuals in their private lives. It is, however, subject to the objection that it recognizes no moral constraints (such as rights) on the way in which the moral good may be promoted. For simplicity, I set this concern aside in what follows.

Adler assumes that the relevant wellbeing for morality is lifetime brute luck wellbeing. I agree, and wellbeing should hereafter be so understood. He further holds that it is only the wellbeing of “persons” that matters morally, where persons are beings that are self-aware, language-using, and capable of deliberation and planning (pp. 4-10). This, I believe, mistakenly excludes the wellbeing of lower but sentient animals, but I won’t pursue the issue here.

2. The Moral Goodness of Outcomes

Adler defends the following theory of moral betterness of outcomes:

Continuous Prioritarianism: O1 is morally at least as good as O2 just in case the sum, over all individuals, of “morally weighted” (brute luck, person-centered) wellbeing in O1 is at least as great as that in O2, where the morally weighted wellbeing is a strictly increasing and strictly concave function of the wellbeing (the same function for all individuals).

Strict concavity in wellbeing just ensures that there is decreasing marginal moral value for increases in wellbeing (e.g., the moral value of increasing wellbeing from 100 to 101 is less than increasing it from 10 to 11). Continuous prioritarianism is just like utilitarianism applied to lifetime, brute luck wellbeing, except that wellbeing has decreasing marginal moral value.

Unlike leximin, it does not give absolute priority to those who are worse off.

Continuous prioritarianism follows from the following conditions:

Anonymity: If one outcome has the same pattern of wellbeing as another, but involves a permutation of wellbeing values among finitely many people, then the two outcomes are equally morally good.

For example (for a two person-case), if X produces distribution $\langle 2,1 \rangle$ and Y produces $\langle 1,2 \rangle$, then the two are equally good. This is incompatible with differential desert or responsibility claims, but Adler rightly rejects the appeal to desert and he assumes that brute luck wellbeing (not outcome wellbeing) is at issue.

Pareto Indifference (wellbeing version): If each person is equally well off in two outcomes, then the two outcomes are equally morally good.

For example, if X produces $\langle 2,1 \rangle$ and Y also produces $\langle 2,1 \rangle$, then X and Y are equally good. This condition rules out the relevance of features independently of their impact on wellbeing.

Strong Pareto Superiority (wellbeing version): If everyone is at least as well off in one outcome as in another, and at least one person is better off, then the first outcome is morally better than the second.

For example, $\langle 2,1 \rangle$ is morally better than $\langle 1,1 \rangle$. This condition requires that moral goodness be

positively sensitive to the promotion of individual wellbeing.

Pigou Dalton: Transferring a fixed amount of wellbeing from a better off person, b , to a worse off person, w , when this does not make w better off than b , makes things better, if no one else is affected.

For example, $\langle 4, 3 \rangle$ is better than $\langle 5, 2 \rangle$. This condition imposes a very weak priority for the worse off and rules out utilitarianism, for example.

Once we set aside rights considerations, the above conditions are, I believe, plausible. The remaining two conditions, I would argue, are not.

Continuity (simplified version): For any given admissible method of measuring the wellbeing of individuals, if outcome x is morally better than outcome y , then there is an outcome, x' , for which everyone's wellbeing is less than in x , and an outcome, y' , for which everyone's wellbeing is greater than in y , such that x' is morally better than y' .

For example, if $\langle 9, 1 \rangle$ is morally better than $\langle 8, 1 \rangle$, then for some sufficiently small positive n , n' , m , and m' , $\langle 9-m, 1-n \rangle$ is better than $\langle 8+m', 1+n' \rangle$. This rules out leximin, since, no matter how small the numbers are, the worst-off person in $\langle 9-m, 1-n \rangle$ will be worse off than the worst-off person in $\langle 8+m', 1+n' \rangle$.

Continuity is a desirable formal property, but I'm skeptical that it is a core requirement of adequacy (although related properties about tradeoffs may be).

For the next condition, let us say that, relative to two outcomes, an individual is

unaffected just in case she has the same level of wellbeing in each outcome.

Separability (of people): The ranking of two outcomes is fully determined by the wellbeing levels of the affected individuals (and does not depend on the level of wellbeing of unaffected individuals).

For example, $\langle 10, 1, 10 \rangle$ is at least as good as $\langle 10, 0, 12 \rangle$ if and only if $\langle 100, 1, 10 \rangle$ is at least as good as $\langle 100, 0, 12 \rangle$ (in both cases, the first person is unaffected). This rules out, for example, sensitivity to issues of equality of wellbeing. A measure of equality might, for example, hold that $\langle 10, 1, 10 \rangle$ is better than $\langle 10, 0, 12 \rangle$ on the ground that the person losing 2 units remains above the average but also hold that $\langle 100, 1, 10 \rangle$ is worse than $\langle 100, 0, 12 \rangle$ on the ground that the person losing 2 units is already below average. I see no reason to rule out sensitivity to equality.

The main objection against continuous prioritarianism is that it has the following property:

Numbers Win Property: A decrease, d , in wellbeing to one (e.g., worst off) person with wellbeing w , combined with *arbitrarily small* increases in wellbeing, as a percentage of d , to each of enough people who are *arbitrarily highly* well off, as multiples of w , makes things morally better.

This entails that the combination of (1) a loss (no matter how large) to the worst off person, combined with (2) a trivially small gain (as a percentage of the above loss) to several others, no

matter how well off they are (as a multiple of the worst off person's wellbeing), makes things better, as long as there are enough others. The number of gainers can trump the priority for the worse off. This seems problematic. Adler acknowledges this, but he argues that all alternative prioritarian views have even more problematic implications.

Leximin, for example, satisfies all the above conditions other than Continuity. It has, however, the opposite problem:

Absolute Priority for the Worse Off: An increase, i , in wellbeing of a given person, combined with *arbitrarily large* decreases in wellbeing, as a percentage of i , to each of an *arbitrarily large* number of people who remain (perhaps only slightly) better off than the first person, makes things morally better.

I agree with Adler that this is more problematic than the Numbers Win Property.

Another possibility (described by Tungodden 2000 and developed by Brown 2005) is that wellbeing has decreasing marginal moral value, and there is also some *threshold* level of wellbeing (e.g., the level for minimally adequate life) such that increases in wellbeing below that level have *absolute priority* over increases in wellbeing above that level. Let the *truncated* prioritarian value of a given level of wellbeing be the prioritarian value of the lesser of that level of wellbeing and the threshold. For example, if the threshold is 16, then the truncated prioritarian value of 12 is just the prioritarian value of 12, whereas the truncated prioritarian value of 25 is the prioritarian value of 16. Consider, then:

Threshold prioritarianism: Maximize total truncated (at threshold) prioritarian total value

first; if there is tie, then maximize the prioritarian total value.

For example, suppose that the prioritarian value of a level of wellbeing is the square root of that level (which produces decreasing marginal value) and the wellbeing threshold is 16. Threshold prioritarianism says that $\langle 4, 4, 16 \rangle$ (total truncated value of $2+2+4$, or 8) is morally better than $\langle 1, 1, 100 \rangle$ (total truncated value of $1+1+4$, or 6), but morally worse than $\langle 0, 16, 100 \rangle$ (total truncated value of $0+4+4$, or 8). The first is worse than the third because they are tied in truncated value (8), but the first has lower total value ($2+2+4$, or 8 versus $0+4+10$, or 14).

Threshold Prioritarianism satisfies all the above properties other than Continuity. It also avoids both of the above problems. Adler points out, however, that it faces the problem of justifying the threshold (life worth living? minimally adequate life? minimally good life?). Moreover, one might object to the fact that it gives absolute priority to a benefit to one person below the threshold over arbitrary large multiples of that benefit to arbitrarily many people who are above the threshold. Still, this seems like a contender, but I won't explore it here.

3. Moral Goodness of Actions

Let us assume that Adler's continuous prioritarianism is the correct theory for assessing the moral goodness of *outcomes*. How is the moral goodness of *actions* determined? If each action produces, with certainty, an outcome, then the consequentialist answer is trivial. Each action is evaluated on the basis of the certain outcome that it produces. In the more general case, however, each action only produces a probability distribution over outcomes. How are actions to be evaluated in that case?

Adler defends:

Ex Post Welfaristic Consequentialism: The moral ranking of actions is determined by the ranking of the *expected moral value* of their distributions of individual wellbeing in their possible outcomes.

This is the standard version of consequentialism. I, however, find the following more plausible:

Ex Ante Welfaristic Consequentialism: The moral ranking of actions is determined by the ranking of the *moral value* of their distributions of expected individual wellbeing in their possible outcomes.

Both approaches appeal to expected values of the actions, but they do so in different ways. To see the difference, suppose that one action produces a 100% chance of a 5-5 distribution, while a second action produces a 50% chance of 12-0 and a 50% chance of 0-12. Suppose that 5-5 is morally better than 12-0 (and 0-12) on prioritarian grounds. Ex post prioritarianism then says that the first action is morally better than the second (since the first is certain to produce a more valuable outcome than the second). Ex ante prioritarianism, by contrast, does *not* base the moral value of actions on the *expected moral value* of the outcomes they produce. Instead, it bases it on the moral value of the *distribution of expected wellbeing* for individuals. The first action has an expected wellbeing of 5 for each individual, whereas the second action has an expected wellbeing of 6 for each. Given that each person has a higher expected wellbeing in the second, the ex ante version judges the second action to be morally better.

As the above example shows, ex post prioritarianism violates the ex ante Weak Pareto

Superiority condition (based on expected wellbeing). In the example, both individuals are ex ante better off with the second action, but ex post prioritarianism judges the first action to be morally better. I take that to be a near-decisive objection to the view, although Adler does an excellent job of defending the view.

I believe that the moral goodness (or permissibility) of an action is fully determined by the facts (including facts about probabilities) at the time of choice. Thus, I believe that ex ante consequentialist theories are more plausible than their ex post counterparts. In particular, ex ante continuous prioritarianism is, I claim, more plausible than the ex post version.

Adler insightfully points out, however, that ex ante Pigou-Dalton is incompatible with three seemingly plausible principles connecting values of actions with the values of their outcomes.

Sure Thing Principle (simplified version): If, for some set, S , of outcomes (1) A and B assign the same probabilities to the outcomes in S , and (2) there is shift of probability from one or more outcomes in S to other outcomes (which need not be in S) such that (a) C is obtained from A by such a shift, and (b) D is obtained from B by the very same shift, then A is at least as good as B if and only C is at least as good as D .

For example, consider the following four actions:

	Wellbeing		
	.5	.5	Expected Wellbeing
A	80-20	50-50	65-35
B	90-10	50-50	70-30

C	80-20	10-200	45-110
D	90-10	10-200	50-105

Here C is like A, and D is like B, except that the .5 probability on 50-50 has been shifted to 10-200. Because A and B only differ with respect to the first column, and the same is true of C and D, the Sure Thing Principle requires that A be ranked in comparison with B in the same way that C is ranked in comparison with D. Ex Ante Pigou Dalton, however, requires that (1) $A > B$ (since, in comparison with B, A shifts expected wellbeing from the better off to worse off) but (2) $C < D$ (since, in comparison with D, C shifts expected wellbeing from the worse off to better off). Ex ante Pigou Dalton seems more plausible to me. It rightly focuses on the distribution of expected wellbeing.

Ex ante Pigou Dalton (in conjunction with Anonymity) is also incompatible, Adler shows, with two further standard principles:

Stochastic Dominance: If one probability distribution can be obtained from another by shifting probabilities toward morally better outcomes, then the former is morally better than the latter.

Weak Stochastic Dominance: If every outcome with non-zero probability under A is morally better than every outcome with non-zero probability under B, then A is better than B.

I will focus on the latter, which is strictly entailed by the former.

Consider two actions, one of which produces a 50% chance of 90-10 and a 50% chance of 10-90, and the other of which produces a 100% chance of 80-20. Because the first produces

expected wellbeing of 50-50 and the second produces expected wellbeing of 80-20, ex ante Pigou Dalton entails that the first is better than the second. Pigou Dalton for outcomes (combined with Anonymity), however, entails that each possible outcome of the first (90-10 and 10-90) is worse than each possible outcome of the second (80-20). Weak Stochastic Dominance thus entails that second is better than the first, which violates ex ante Pigou Dalton. Again, the latter seems more plausible to me.

The implausibility of Weak Stochastic Dominance for prioritarian moral goodness can, I claim, be further illustrated by noting that, combined with Pigou Dalton for outcomes, it also violates ex ante Pareto Indifference and ex ante Strong Pareto Superiority. Consider, for example, an action with 100% chance of 5-5 and an action with a 50% chance of 6-4 and a 50% chance of 4-6. Ex ante Pareto Indifference requires that the two be judged equally morally good, whereas Weak Stochastic Dominance combined with ex post Pigou Dalton requires that the first be judged morally better. Moreover, assuming Continuity, for a sufficiently small positive number, ϵ , these two principle also judge $(5-\epsilon)-(5-\epsilon)$ as morally better, and that violates ex ante Strong Pareto Superiority. Again, the ex ante principles seem much more plausible to me. To his credit, Adler insightfully discusses these very points.

The Sure Thing Principle, Stochastic Dominance, and Weak Stochastic Dominance are each highly plausible principles when applied to standard accounts of *prudential value*. They are also plausible principles, if any prioritarianism (or risk aversion) present in moral value is also present in prudential value, and vice-versa. For example, there is no conflict (1) when utilitarian moral value is combined with standard prudential value (neither of which involve any prioritarianism), or (2) when prioritarian moral value is combined with the same prioritarianism applied to prudential value (which is fairly implausible). In that case, the above three principles

do not conflict with the ex ante principles. These principles are, however, implausible, I claim, when moral value is prioritarian and prudential value is not.

In short, I claim that that the moral assessment of actions does not supervene on the probabilities and moral goodness of the outcomes they might produce. It supervenes instead on the moral value of the expected wellbeing for individuals. Thus, the ex ante version of continuous prioritarianism is, I claim, more plausible than the ex post version. Of course, this is all highly contested, and Adler does a wonderful job of explaining the issues and defending his position.

In summary, Adler's book is state-of-the art, and all the objections that I raised above (and many others!) are anticipated and addressed in depth. Because many of the issues are technical and complex, the book is not an easy read, but the book does a superb job at presenting issues in highly intuitive terms. Anyone with serious interests in consequentialist moral theory should read this book.

Peter Vallentyne

University of Missouri

Columbia, MO 65211, USA

Vallentynep@missouri.edu

References

Brown, C. 2005. Priority or Sufficiency ... or Both. *Economics and Philosophy* 21:199-220.

Tungodden, B. 2000. Egalitarianism: Is Leximin the Only Option? *Economics and Philosophy* 16: 229-246.