

The Separateness of Persons Objection

Table of Contents

1. Introduction
2. Compensation
3. Fungibility
4. Anti-Aggregative Intuitions
 - Reinterpreting the Intuition
 - Debunking the Intuition
 - Rivals Fare No Better
5. Conclusion
6. Resources and Further Reading

Introduction

It's widely agreed that self-interest (or prudence) calls for aggregating harms and benefits across different moments *within* one's life, to maximize one's overall well-being. For example, visiting the dentist is prudent despite the immediate unpleasantness because it helps to avert greater harm to one's future self. Aggregative consequentialist theories like utilitarianism go one step further: they aggregate harms and benefits *between* different people's lives to maximize overall societal well-being. It can be worth imposing harms on some individuals, utilitarians claim, if that prevents greater harms to others. This leads some critics to claim that utilitarianism neglects the moral significance of the boundaries between individuals.

This *separateness of persons* objection was stated most famously by John Rawls:

[Utilitarianism] is the consequence of extending to society the principle of choice for one man, and then, to make this extension work, conflating all persons into one... Utilitarianism does not take seriously the distinction between persons. ¹

Despite its influence, the reasoning behind this objection can be difficult to pin down. The idea that utilitarians must "conflat[e] all persons into one" seems to presume that they (i) *start* with "the principle of choice for one man", and then (ii) argue for their view on the basis that all of society can be treated (perhaps metaphorically) as just another individual. But in fact many utilitarian arguments, as laid out in Chapter 3: [arguments for utilitarianism](#), do *not* take this form. So this

interpretation of the objection seems too narrow. It might debunk one particular argument for utilitarianism, but utilitarianism itself may still be well-supported on other grounds.

On a broader interpretation, we may take the objection to assert that respecting the distinction between persons requires treating inter-personal tradeoffs (those *between* lives) differently from intra-personal ones (those *within* one life). On this interpretation, it's the implications of utilitarianism, not what argument led to it, that are seen as objectionable. But what is the positive case for treating intra- and inter-personal tradeoffs differently?² The remainder of this article explores three candidate arguments, based on (i) compensation, (ii) fungibility, and (iii) anti-aggregative intuitions.

Compensation

The standard interpretation of the separateness of persons understands it to be a matter of *compensation*.³ The agent who gets harmed is *compensated* if they later receive a greater benefit as a result, whereas they receive no such compensation if the benefit goes to someone else. As Nozick put it: "To use a person [for another's benefit]... does not sufficiently respect and take account of the fact that he is a separate person, that his is the only life he has. *He* does not get some overbalancing good from his sacrifice".⁴

Utilitarians may respond that every person in need is an individual, equally deserving of moral concern and respect, living the only life *they* have.⁵ To make the challenge vivid, we may imagine ourselves in the position of the one being used as a means. When utilitarianism demands that we bear a cost, it does not deny that the beneficiaries are distinct from us. It merely denies the egoistic assumption that their distinctness means that they should not matter to us; it denies that we could reasonably demand a *veto* over every trade-off in which our interests are negatively implicated.⁶ And, indeed, this is hardly unique to utilitarians. As Katarzyna de Lazari-Radek and Peter Singer write, "Anyone who supports taxing people on high incomes and using the revenue to provide benefits to others in need must agree that it is sometimes justifiable to impose a cost on one person to benefit another."⁷

Fungibility

A deeper concern is that utilitarianism might seem to treat individuals as entirely fungible or replaceable without regret. As Peter Singer characterizes the view: "It is as if sentient beings are receptacles of something valuable, and it does not matter if a receptacle gets broken so long as there is another receptacle to which its contents can be transferred without any getting spilt."⁸

To make the problem vivid, imagine that Connie the Consequentialist is faced with two poison victims, and just enough anti-venom to save one of them.⁹ And suppose that, faced with their pleading faces, but realizing that it makes no difference to the total amount of well-being which

person she saves, Connie finds herself feeling completely *indifferent* about her choice. It's as if she had to choose between a \$20 bill or two tens.

It seems that Connie is making a deep moral mistake here. She is treating the two people's interests as completely fungible, like money, and neglecting the fact that each person is of distinct intrinsic importance, in their own right, and not merely a fungible means to aggregate well-being. The correct moral theory, we feel, must attribute intrinsic value to particular individuals and not just to the sum total of their well-being.¹⁰

One of the authors (Chappell) has argued that there is no barrier to utilitarianism assigning intrinsic value to individuals in this way:¹¹

There is not just one thing, the global happiness, that is good. Instead, there is my happiness, your happiness, Bob's, and Sally's, which are all equally weighty but nonetheless distinct intrinsic goods. What this means is that the morally fitting agent should have a corresponding plurality of non-instrumental desires: for my welfare, yours, Bob's, and Sally's. Tradeoffs between us may be made, but they are acknowledged as genuine tradeoffs: though a benefit to one may *outweigh* a smaller harm to another, this does not *cancel* it. The harm remains regrettable, for that person's sake, even if we ultimately have most reason to accept it for the sake of more greatly benefiting another.

On this view, what it's appropriate for Connie to feel is not *indifference*, but rather *ambivalence*—to be genuinely torn, as she is pulled (equally strongly) in different directions by the conflicting interests of the two individuals who need her help. In this way, the utilitarian can avoid treating individuals as fungible, and instead fully recognize and appreciate each individual's separate value.

This utilitarian response rejects the assumption that *commensurability* of value (that is, the ability to compare and make tradeoffs between competing interests or values) entails *fungibility* of value (or replaceability without regret). The assumption may be rejected since non-fungibility can be secured by having multiple genuinely *distinct* values, which may nonetheless be weighed against each other.¹²

This is perhaps clearest when considering other forms of (putative) value: an art-lover might intrinsically value each item in their art collection, and yet still be willing to smother a fire with one painting if that would save five others. In valuing each painting separately, they will regret the loss of the unique painting that is thereby destroyed. But they may regard the sacrifice as worth it, even so. Nothing in their attitude here betrays an objectionably instrumental attitude towards their artworks. Just as this art consequentialist respects the separateness of paintings, so the utilitarian (or welfare consequentialist) respects the separateness of persons.

Anti-Aggregative Intuitions

Finally, critics might grant that utilitarianism really does separately value individual persons, but just *not in the right way*. On this view, there is something intuitively problematic about utilitarian [aggregation](#). Instead of summing different people's interests, strictly anti-aggregative views might advocate for a *maximin* approach that simply seeks to improve the position of the worst-off. As Nagel writes:

Where there is a conflict of interests, no result can be completely acceptable to everyone. But it is at least possible to assess each result from each point of view to find the result that is the least unacceptable to the person to whom it is most unacceptable... A radically egalitarian policy of giving absolute priority to the worst-off, regardless of numbers, would result from always choosing the least unacceptable alternative, in this sense.¹³

Maximin has severe problems as an alternative to utilitarian aggregation. Taking maximin literally, it would be preferable to give a lollipop to the most miserable person on Earth rather than, say, to avert a global pandemic or nuclear war that would happen after the most miserable person's death (or that would somehow leave this individual unharmed).

Still, there are particular cases in which utilitarian aggregation seems (intuitively) to yield the wrong result. Consider Scanlon's famous *Transmitter Room* case:¹⁴

Jones has suffered an accident in the transmitter room of a television station. To save Jones from an hour of severe pain, we would have to cancel part of the broadcast of a football game, which is giving pleasure to very many people.¹⁵

Intuitively, it does not matter how many people are watching the football game; it's simply *more important* to save Jones from suffering severe pain during this time.¹⁶

Reinterpreting the Intuition

Why is it more important to save Jones? One answer would be that we cannot aggregate distinct interests, so all that is left to do is to satisfy whichever individual moral claim is strongest, namely, Jones's. But Parfit suggests an alternative—[prioritarian](#)—explanation: perhaps we should help Jones because he is much worse off, and thus has greater moral priority.¹⁷

(While utilitarians reject this prioritarian claim, they may nonetheless take comfort if it turns out that our intuitions are more closely aligned with prioritarianism than with anti-aggregationism. This is for two main reasons. First, they may regard prioritarian intuitions as easily debunkable. And second, they may regard prioritarianism as [close enough](#) to utilitarianism that they are not so concerned to press the dispute.)

Parfit argues that his prioritarian account is preferable to Scanlon's anti-aggregative approach in cases where the two diverge. We can see this by imagining cases in which the many smaller benefits would go to some of the worst-off individuals. For example, it would clearly be better to give an

extra five years of life to each of a million child cancer victims, than to give fifty more years of life to a single adult. In rejecting aggregation, we might have to prioritize a single large benefit to someone already well-off, rather than (individually smaller but collectively immensely larger) benefits to a great many worse-off individuals. That seems clearly wrong. It would not, for example, be a good thing to take a dollar from each of a billion poor people in order to give a billion dollars to someone who was wealthy to begin with.

So, rather than refusing to aggregate smaller benefits, Parfit suggests that we should simply weigh harms and benefits in a way that prioritizes the worse-off. Two appealing implications of this view are that: (i) We generally should not allow huge harms to befall a single person, if that leaves them much worse off than the others with competing interests. (ii) But we should allow (sufficient) small benefits to the worse-off to (in sum) outweigh a single large benefit to someone better-off. Since we need aggregation in order to secure claim (ii), and we can secure claim (i) without having to reject aggregation, it looks like our intuitions are overall best served by accepting an aggregative moral theory.

Debunking the Intuition

Common intuitions suffer from *scope insensitivity*, reflecting our inability to truly grasp large numbers.¹⁸ Our intuitions do not respond very differently to whether the number of competing interests is a million, a billion, or a [googolplex](#). But the real difference in value between these numbers is immense. So we should not trust our intuitions when they treat these vastly different numbers as morally alike. Utilitarians may thus feel comfortable rejecting anti-aggregative intuitions as especially untrustworthy.

Despite Parfit's arguments described above, even prioritarrians may ultimately need to follow the utilitarian in accepting a debunking response. Consider: critics may insist that Parfit's prioritarian account cannot do full justice to our starting intuition about the Transmitter Room case. Granted, sufficient priority weighting may explain how Jones's suffering can outweigh the aggregate pleasure of a million, or even a billion, better-off football fans. But so long as the priority weighting is finite, there will be *some* (perhaps astronomically large) number of smaller pleasures that could, in theory, outweigh Jones's suffering. At this point, defenders of aggregation may simply accept this implication, and suggest that any residual intuitive discomfort with this conclusion is best explained as a mistake resulting from scope neglect.

Rivals Fare No Better

Contrary to utilitarianism, one might be tempted to think that some benefits are so *trivial* that we should round them down to zero, rather than allowing vast numbers of them to sum to something morally significant. But Parfit proves that this way of thinking is a mistake. To see why, consider the following plausible-seeming claim:

(P): We ought to give one person one more year of life rather than lengthening any number of other people's lives by only one minute.¹⁹

One year is about half a million minutes. So Parfit invites us to imagine a community of just over a million people, and apply the choice described in (P) to each of them. Each person in the community would then gain one year of life. But consider the opportunity cost. If each time we had instead given one more minute of life to everyone else, the end result would be a gain of *two* years of life for each person. So the choice described in (P), when repeated in this way, results in everyone being worse off than they otherwise would have been.

This clearly shows that (P) is a bad principle *in iterative contexts* like that described above. Does it show that (P) is a bad principle even in a one-off application? That is less immediately clear, but we may be able to show this with further argument.²⁰ Parfit himself appeals to a distinction between fundamental moral *principles* and mere *policies* (or rules of thumb), suggesting that only the latter should be contingent on context in this way. If he is right about that, this would suggest that our fundamental moral principles must allow for unrestricted aggregation, in contrast to claims like (P). Any rule that we take to apply only some of the time (e.g., in one-off applications but not iterated ones) must, for Parfit, be a mere rule of thumb rather than a fundamental moral principle.

We might supplement Parfit's argument by observing that the expected value of each choice described above is *independent* of the other choices being made. The value of giving everyone one more minute (just once) is the same as the value of giving everyone one more minute (for the millionth time).²¹

This is important for two reasons. First, independence implies that the expected value of the one-off choice is equal to the average value of the repeated choice. So, since repeatedly choosing *one minute for everyone* is more worthwhile than repeatedly choosing *one year for one person*, it follows (from independence) that the former choice is also more worthwhile in the one-off case. This is a surprising and important result.

The second reason why the independence claim is important here is that it can help to shed light on why this initially surprising result makes sense, and is plausible upon reflection. Whatever valuable events an extra year of life would offer—all the moments of happiness, completed projects that would otherwise have been cut short, etc.—we should expect twice as many such events to be enabled by offering an extra minute of (representative) life to each of a million people.

Putting this all together, then, defenders of aggregationism can offer a two-pronged response to their critics. First, our initial anti-aggregative intuitions may be explained away. And second, further reflection shows that anti-aggregative principles would have implications that are arguably more objectionable than those of aggregationism.

Conclusion

We've seen that the separateness of persons objection to utilitarianism can take three forms, none of which is decisive. The *compensation* objection rests on implausibly egoistic assumptions. The *fungibility* objection involves a misconception: utilitarianism need not treat distinct individuals as fungible. Finally, while *anti-aggregative* intuitions have some force, we've also seen that there are ways for utilitarians to resist them, and that these anti-aggregative views face even worse difficulties.

Next: The Demandingness Objection

Other Objections to Utilitarianism

How to Cite This Page

```
Chappell, R.Y. and Meissner, D. (2023). The Separateness of Persons Objection. In R.Y. Chappell, D. Meissner, and W. MacAskill (eds.), An Introduction to Utilitarianism, <https://www.utilitarianism.net/objections-to-utilitarianism/separateness-of-persons>, accessed 12/3/2024.
```

Resources and Further Reading

- David Brink (2020). [Consequentialism, the Separateness of Persons, and Aggregation](#). In Douglas W. Portmore (ed.) *The Oxford Handbook of Consequentialism*. Oxford University Press.
- Richard Y. Chappell (2021). [Parfit's Ethics](#), section 3.2. Cambridge University Press.
- Richard Y. Chappell (2015). [Value Receptacles](#). *Noûs*, 49(2): 322–332.
- G.A. Cohen (2011). [Rescuing Conservatism: A Defense of Existing Value](#). In R. Jay Wallace, Rahul Kumar, and Samuel Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*. Oxford University Press.
- Stephan Dickert, Daniel Västfjäll, Janet Kleber, and Paul Slovic (2015). [Scope insensitivity: The limits of intuitive valuation of human lives in public policy](#). *Journal of Applied Research in Memory and Cognition*, 4(3): 248–255.
- Robert Nozick (1974). *Anarchy, State, and Utopia*. Basic Books.
- Derek Parfit (2003). [Justifiability to Each Person](#). *Ratio*, 16(4): 368–390
- Ellen Frankel Paul, Fred D. Miller, Jr., and Jeffrey Paul (eds.) (2010). [Utilitarianism: The Aggregation Question](#). Cambridge University Press.
- John Rawls (1971). *A Theory of Justice*. Belknap Press.
- T.M. Scanlon (1998). *What We Owe to Each Other*. Belknap Press.

-
1. Rawls, J. (1971). *A Theory of Justice*, p.27. ↩
 2. As [R.M. Hare](#) presses the challenge: “To have concern for someone is to seek his good, or to seek to promote his interests; and to have equal concern for all people is to seek equally their good, or to give equal weight to their interests, which is exactly what utilitarianism requires. To do this is to treat others’ interests in the same way as a prudent person treats his own interests, present and future... To do this is not to fail to ‘insist on the separateness of persons’.” See: Hare, R.M. (1985). ‘Rights, Utility, and Universalization: Reply to J.L. Mackie’, in R. Frey (ed.) *Utility and Rights*, p.107. ↩
 3. Brink, D. (2020). [Consequentialism, the Separateness of Persons, and Aggregation](#). In Douglas W. Portmore (ed.) *The Oxford Handbook of Consequentialism*. Oxford University Press. ↩
 4. Nozick, R. (1974). *Anarchy, State, and Utopia*, p.33. ↩
 5. An interestingly different line of response would be to appeal to a [Veil of Ignorance](#) argument. While the individual who gets harmed is not compensated for it *in the moment*, each individual should be willing *in advance* (i.e., from behind the veil of ignorance) to agree to utilitarian tradeoffs since this is the best way for them to maximize their own well-being, in expectation. ↩
 6. To expand on this point: If we reject egoism, we should also reject the assumption that compensation is essential for moral justification. For on any other view, we may be morally *required* to accept some costs for others’ sakes. Next note that we need not be granted a veto power that it would not be morally legitimate for us to exercise. So, some *may* use us, for another’s benefit, compatibly with full moral respect, if the use is one that we are morally required to accept.

Brink notes that Rawls’s own egalitarian view often requires uncompensated sacrifices from the wealthy. Brink, D. (2020). [Consequentialism, the Separateness of Persons, and Aggregation](#). In Douglas W. Portmore (ed.) *The Oxford Handbook of Consequentialism*. Oxford University Press, p.387. ↩
 7. de Lazari-Radek, K. and Singer, P. (2017). [Utilitarianism: A Very Short Introduction](#). Oxford University Press. Chapter 4: Objections, p. 82. ↩
 8. Singer, P. (2011). *Practical Ethics*, third edition. Cambridge University Press, p. 106 ↩
 9. The following paragraphs draw directly from Chappell, R.Y. (2015). [Value Receptacles](#). *Noûs*, 49(2): 322–332. ↩
 10. Cf. Cohen, G.A. (2011). [Rescuing Conservatism: A Defense of Existing Value](#). In R. Jay Wallace, Rahul Kumar, and Samuel Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy*

of T.M. Scanlon. Oxford University Press. ↩

11. Chappell, R.Y. (2015). [Value Receptacles](#). *Noûs*, 49(2): 322–332, p. 328. ↩
12. “Value pluralism” is often used to refer to the idea of multiple distinct *types* or kinds of values. Instead, the relevant form of pluralism to secure non-fungibility is *token* pluralism. Bob’s and Sally’s interests may both be values of the same *kind* (namely, welfare value), but they are distinct individual (or “token”) values in the sense that it is fitting to have a *separate* intrinsic desire for each. Contrast this with money, where distinct \$20 bills are *not* distinctly valuable: it would be strange to desire each bill separately, rather than just having a single overarching desire for “more money” that any \$20 bill could equally well satisfy. ↩
13. Nagel, T. (1977). Equality. Reprinted in (1979) *Mortal Questions*. Cambridge University Press, p. 123. ↩
14. Scanlon, T.M. (1998). *What We Owe to Each Other*. Belknap Press. ↩
15. This is a quote from Derek Parfit’s concise summary of the thought experiment, on p. 375 of Parfit, D. (2003). [Justifiability to Each Person](#). *Ratio*, 16(4): 368–390. ↩
16. The following subsections draw directly from section 3.2 of Chappell, R.Y. (2021). *Parfit’s Ethics*. Cambridge University Press. ↩
17. Parfit, D. (2003). [Justifiability to Each Person](#). *Ratio*, 16(4): 368–390.

What if the watchers are somehow even worse-off? Then Parfit’s explanation fails, but he could comfortably suggest in this case that slightly improving the lot of billions of worse-off individuals really should be prioritized over offering great relief to a single individual who is already in a better state than these others. ↩

18. Dickert, S., Västfjäll, D., Kleber, J., & Slovic, P. (2015). [Scope insensitivity: The limits of intuitive valuation of human lives in public policy](#). *Journal of Applied Research in Memory and Cognition*, 4(3): 248–255. ↩
19. Parfit, D. (2003). [Justifiability to Each Person](#). *Ratio*, 16(4): 368–390, p.385. ↩
20. This text continues to draw directly from section 3.2 of Chappell, R.Y. (2021). *Parfit’s Ethics*. Cambridge University Press. ↩
21. There are ways of imagining the case where this would not be so. For example, if we imagine giving the extra minutes of life to each person on their deathbed, the first several such minutes might be disproportionately lacking in value, compared to a more representative minute of life. To properly test principles of aggregation, we should imagine a setup where the independence assumption holds—for instance, by supposing that the extra minutes are given to people at some earlier point in their lives, before mortal illness strikes. This makes it clearer

how a single minute might, in some cases, have significant value, by being just what the recipient needed to complete some important life project. ↩