REPLY

Superiority Discounting Implies the Preposterous Conclusion

Mitchell Barrington (iD)

University of Southern California, Los Angeles, USA
Dianoia Institute of Philosophy, Australian Catholic University, Melbourne, Australia
E-mail: barringt@usc.edu

Abstract

Many population axiologies avoid the Repugnant Conclusion (RC) by endorsing Superiority: some number of great lives is better than any number of mediocre lives. But as Nebel shows, RC follows (given plausible auxiliary assumptions) from the Intrapersonal Repugnant Conclusion (IRC): a guaranteed mediocre life is better than a sufficiently small probability of a great life. This result is concerning because IRC is plausible. Recently, Kosonen has argued that IRC can be true while RC is false if small probabilities are *discounted* to zero. This article details the unique problems created by combining Superiority with discounting. The resultant view, *Superiority Discounting*, avoids the Repugnant Conclusion only at the cost of the Preposterous Conclusion: near-certain hell for arbitrarily many people is better than near-certain heaven for arbitrarily many people.

1 The Repugnant Conclusion

The Repugnant Conclusion (RC) is that for any number of people living great lives, it would be better to have a sufficiently greater number of people with lives that are barely worth living (Parfit 1986). Nebel (2019) shows that the Repugnant Conclusion follows from the Intrapersonal Repugnant Conclusion (IRC), alongside plausible auxiliary assumptions. This result is surprising because RC is intuitively false, but IRC is not.

The Intrapersonal Repugnant Conclusion is that for each person that might exist, it is better for them to have a guaranteed mediocre life than a sufficiently small chance of a great life and otherwise nothing. Suppose prospect Z guarantees every person who might exist a mediocre life, while prospect A guarantees those in a small subset great lives and the others nothing. According to IRC, if each person's probability of being in the lucky subset is small enough, they are better off with Z's guaranteed mediocre life than A's gamble on a great life.

The crucial premise in the move from IRC to RC is:

Weak Pareto for Equal Risk: For any egalitarian prospects X and Y, if X is better than Y for each person who might exist in either prospect, then X is better than Y (Nebel 2019: 320).[1]

---

[1] The principle specifically covers egalitarian prospects (those in which everyone who exists is equally well off and every person who might exist has an equal probability of existing) because prospects that are better for everyone might still be worse if they increase inequality.

Weak Pareto gets at the intuitive idea that if some prospect is better for everyone, then it is simply better. And since prospect Z (where everyone is guaranteed mediocre lives) is better than A (where each member of a small subset receives a great life) for every person who might exist, it follows via Weak Pareto that Z is better than A. And since Z is better than A, the certain outcome of Z (a greater number of people living mediocre lives) must be better than the certain outcome of A (a smaller number of people living great lives).[2] And this is the Repugnant Conclusion.[3]

2 Superiority discounting

We can avoid the Repugnant Conclusion by endorsing the *superiority* of great lives over mediocre lives:[4]

> *Superiority*: x is superior to y just in case some quantity of x is better than any quantity of y.[5]

In population ethics, this position is most notably adopted by perfectionists, according to whom, "even if some change brings a great net benefit to those who are affected, it is a change for the worse if it involves the loss of one of the best things in life" (Parfit 1986: 163). Accordingly, a population consisting of any number of great lives (those containing the best things in life) is better than a population consisting of any number of mediocre lives (those without the best things in life); RC is false.

But to resist RC in light of Nebel's argument, advocates of Superiority must reject either IRC or Weak Pareto. Nebel supposes that a perfectionist might reject IRC by claiming that "even if some prospect would, in expectation, bring a net benefit to a person, it is worse for her if it lowers her probability of enjoying the best things in life" (2019: 324). This perfectionist would not accept IRC because *any* chance of a great life is better than none. Ultimately, Nebel rejects this kind of perfectionism because it results in an "absurdly reckless" decision theory that instructs agents to "prefer prospects that will almost certainly be worse for us in pursuit of arbitrarily small chances of enjoying the best things in life" (2019: 324).

Kosonen (2021) suggests Superiority advocates instead deny Weak Pareto by adopting *discounting*:

> *Discounting*: Agents should discount small probabilities to zero.[6]

---

[2]Nebel covers this inference with the uncontroversial principle, Certainty Equivalence, which states that the certain outcome of prospect X is better than the certain outcome of Y if and only if X is better than Y (Nebel 2019: 322).

[3]To be precise, the Repugnant Conclusion is that Z is better than another prospect A⋆. A⋆ is identical to A except that it guarantees a specified subset of people a great life, rather than giving every person a small chance of getting a great life. Since A and A⋆ are equally good, and Z is better than A, Z is also better than A⋆.

[4]Superiority appears in many areas of value theory. Prominent endorsements of Superiority include: Hutcheson (1755: 118); Ross (1930: 150); Glover (1977: 710); Edwards (1979: 69–72); Griffin (1986: 85–86); Crisp (1992: 151); Lemos (1993); Mill (1998: 56); Skorupski (1999: 94–101); Brentano (2009: 106).

[5]This is *Weak Superiority* (Arrhenius and Rabinowicz 2015). *Strong Superiority* is the view that *any* quantity of x is better than any quantity of y.

[6]Proposals that endorse discounting include: Bernoulli (1738); d'Alembert (1761); Borel (1962); Buffon (1777); Condorcet (1785); Kagan (1989: 89–92); Jordan (1994); Aboodi et al. (2008); Hawley (2008); Haque (2012); Buchak (2013); Smith (2014, 2016); Bjorndahl et al. (2017); Chalmers (2017); Lazar (2017); Schwitzgebel (2017); Lee-Stronach (2018); Robert (2018); Tarsney (2018); Monton (2019).

This *Superiority Discounter* (SD) will not chase arbitrarily small probabilities at any cost because they will ignore sufficiently small probabilities entirely. SD will endorse IRC because they will ignore the (sufficiently small) probability of any individual getting a great life under A, and judge that it is better for them to receive a guaranteed mediocre life under Z than nothing. However, even though Z is better than A for every individual, SD will not conclude that Z is (impartially) better than A because the probability that *someone* will acquire a great life is too large to ignore. (Indeed, it is certain.) And as per Superiority, a smaller number of people living great lives is better than a larger number of people living mediocre lives. So, Z is better for every person who might exist, but A is nevertheless better than Z; Weak Pareto is false.

3 The Preposterous Conclusion

While discounting steers superiority theories clear of the Repugnant Conclusion, it guides them to a more uncomfortable result. The Preposterous Conclusion is that, for an arbitrarily great number of people, near-certain hell is better than near-certain heaven (where hell is an arbitrarily horrible life and heaven is an arbitrarily great life). In the remainder of this article, I will demonstrate how Superiority Discounting implies the Preposterous Conclusion. (In what follows, I use "prefer" to mean *judges to be better for the subject in question*.)

I will begin with Kosonen's paradigm case, in which SD prefers a non-negligible chance of living a great life and otherwise nothing (prospect A) to a guaranteed mediocre life (prospect Z). Then, I will alter the decision problem in ways that sweeten Z and dampen A but do not affect this preference ordering. Each alteration will make SD's continued preference for A more preposterous. The end result will be the Preposterous Conclusion.

3.1 The paradigm case

The central commitment of Superiority Discounting is:

> *Risky Non-Repugnance*: q chance (or greater) of obtaining at least one life at a high welfare level a is better than certainty of obtaining any number of lives at a low welfare level z, where q is the smallest probability that should not be discounted down to zero (Kosonen 2021: 212).

Supposing that the value of q is one in one million, this comparison is represented in Table 1. Prospect A offers a one-in-one-million chance of a great life and otherwise nothing, while Z guarantees a mediocre life.

Since both states' probabilities are at least equal to one in one million, neither is discounted. And since A's probability of providing a great life is higher than Z's, SD will prefer A.

3.2 Discounting insurance

In Table 2, we will shift some probability from State 1 into some new state, State 3, whose probability will be just low enough to ignore: $1/1m - \varepsilon$ (where ε is an arbitrarily small number). In State 3, A will result in nonexistence, while Z will result in a great life. So, now both prospects have a small chance of producing a great life. However, Z offers the insurance of at least a mediocre life at the cost of an arbitrarily small decrease in the probability of securing a great life.

Table 1. The paradigm case

	State 1 (1 – 1/1m)	State 2 (1/1m)
A		a
Z	z	z

Table 2. Discounting insurance

	State 1 (1 – 1/1m – (1/1m – ε))	State 2 (1/1m)	State 3 (1/1m – ε)
A		a	
Z	z	z	*a*

SD prefers A to Z because they will ignore State 3 and prefer A due to its higher probability of producing a great life.[7] But this preference seems wrong: Z's guarantee of a mediocre life seems well worth the price of an arbitrarily small decrease in the probability of getting a great life. It is not that SD has become absurdly reckless again: they are not chasing an arbitrarily small probability at any cost; they are chasing an arbitrarily small *increase* in probability at any cost. We invoked discounting to escape the former only to end up with the latter, and it is not clear that this result is any better.[8]

3.3 Superiority insurance

The discounting insurance problem arose because discounting makes agents insensitive to insufficiently probable states. But a commitment to Superiority also harbors an insurance problem because it makes agents insensitive to insufficiently *valuable* outcomes. Consequently, SD will be unwilling to buy insurance on their gamble on a great life at an arbitrarily small price to the value of that life.

---

[7]We could instead cash out the discounting strategy as applying to *differences* in prospects' probabilities of producing a great life, rather than the absolute probabilities of states. So, for any prospects X and Y, if X's probability of producing a great life is sufficiently less than Y's, then Y is better than X. Unfortunately, this proposal leads to transitivity violations: if A's probability (of producing a great life) is not sufficiently greater than B's, B might be better than A; if B's probability is not sufficiently greater than C's, then C might be better than B. Nevertheless, A's probability might be sufficiently greater than C's such that A *must* be better than C. So, even though C > B and B > A, A > C.

[8]SD might respond that the strategy of inserting a state *slightly* below the threshold will not work if the threshold is vague, because whether its probability is below the threshold will be indeterminate. However (assuming truths of classical logic are determinate, and determinacy is closed under logical entailment), this will not pose a problem for the argument. SD is right that for no pair of numbers an arbitrarily small distance apart n and $n - ε$ is it the case that determinately, n is not discounted but $n - ε$ is. Nevertheless, it *is* determinate that *there exists* some pair of numbers an arbitrarily small distance apart n and $n - ε$ such that n is not discounted but $n - ε$ is. And so determinately, a state whose probability is an arbitrarily small amount less than the threshold (wherever it is) will not affect SD's preferences. (By analogy, approaches to vagueness that endorse classical logic, including epistemicism and supervaluationism, maintain that determinately, there is a pair of numbers n and $n + 1$ such that with n hairs on one's head one is bald, but with $n + 1$ hairs one is not bald, while conceding that for no pair of numbers n and $n + 1$ is it the case that determinately, with n hairs one is bald, but with $n + 1$ hairs one is not bald.) This point applies *mutatis mutandis* to the threshold separating superior objects from inferior objects, discussed in section 3.3.

SD maximizes their probability of living a great life. So, except in the event of a tie, they will be insensitive to all outcomes in which they do not acquire a great life – even if the life they acquire is arbitrarily close to the greatness threshold.[9] As a result, increasing Z's insurance payout to any value below the threshold will not alter SD's preference. For illustration, suppose that great lives are those at or above welfare level 80. In Table 3, we have increased Z's insurance payout to a life at welfare level 80 – ε. (Call lives that are slightly below the greatness threshold *very good* lives.)

SD continues to prefer A to Z, despite the high probability of a very good life under Z because the probability of acquiring a *great* life remains higher under A. But now the preference of A looks very strange indeed. Given that A and Z have an approximately equal probability of producing the same great life, they should, on this basis alone, be approximately equally good. Add the fact that Z also *guarantees* a life valued approximately equally to the one A takes a one-in-a-million gamble on, and Z should be much better (on the order of one million times better, if such quantifications make sense) than A.

3.4 Extreme values in excluded states

In 3.2, we introduced State 3 to establish the possibility of Z producing a great life. Since SD will ignore it due to its improbability, we can make the outcomes in this state arbitrarily extreme without affecting their preference. Suppose we raise the welfare level of the life Z would produce in State 3 from 80 to 100 and lower that of A's life to −100. (Call lives at welfare level 100 *heaven* and those at welfare level −100 *hell*.)

In Table 4, A still offers a one-in-a-million chance of a life at welfare level 80, but it now offers approximately the same chance of hell (and otherwise nothing). Z again guarantees a very good life, but it now offers a small chance of heaven.

SD still prefers A because they will continue to ignore State 3 due to its improbability. But intuitively, A is now worse than certain nonexistence: the small probability of a great life does not seem worth the risk of an approximately equal probability of hell. On the other hand, Z *guarantees* a very good life alongside a small chance of heaven; it is clearly (and considerably) better than certain nonexistence. Nevertheless, SD prefers A not only to certain nonexistence, but also to Z.

3.5 Ignoring negative value

Since SD maximizes the probability of acquiring a great life, their preferences are only responsive to outcomes in which they acquire a great life (except in the event of a tie). In 3.3, this feature allowed us to increase Z's insurance payout to a welfare level of just under 80 without impacting their preference. But there is another problem we can use this feature to exploit. SD does not merely ignore lives marginally below the greatness threshold; they also ignore the possibility of procuring a life that is worse than nonexistence.

In Table 5, A produces a life at −100 in State 1. As in 3.3, this problem is not the result of discounting insufficiently probable states (indeed, State 1 is the *most* probable

---

[9]SD might suppose that there is no *absolute* threshold: for all welfare levels *x* and *y*, if *x* is sufficiently (e.g., 25 points) greater than *y*, then lives at *x* are superior to lives at *y*. But this, too, violates transitivity. For any number of lives at welfare level 80, there is some number of lives at 60 that would be better; and for any number of lives at 60, there is some number of lives at 40 that would be better. By transitivity, for any number of lives at 80, there is some number of lives at 40 that would be better. But this contradicts the stipulation that lives at 80 are superior to lives at 40 (in virtue of the greater-than-25-point difference in welfare levels).

Table 3. Superiority insurance

	State 1 $(1 - 1/1m - (1/1m - \varepsilon))$	State 2 $(1/1m)$	State 3 $(1/1m - \varepsilon)$
A		80	
Z	**80 – ε**	**80 – ε**	80

Table 4. Extreme values in excluded states

	State 1 $(1 - 1/1m - (1/1m - \varepsilon))$	State 2 $(1/1m)$	State 3 $(1/1m - \varepsilon)$
A		80	**−100**
Z	80 – ε	80 – ε	**100**

Table 5. Ignoring negative value

	State 1 $(1 - 1/1m - (1/1m - \varepsilon)$	State 2 $(1/1m)$	State 3 $(1/1m - \varepsilon)$
A	**−100**	80	**−100**
Z	80 – ε	80 – ε	100

state) but of ignoring insufficiently valuable outcomes. SD's focus on maximizing their probability of acquiring a great life makes them insensitive to this extremely bad, extremely probable outcome.

SD continues to prefer A because, in maximizing their probability of living a great life, they ignore the high probability of hell under A. Unsurprisingly, ignoring very bad, very probable outcomes has unfortunate results. A now offers a one-in-a-million chance of *not* ending up in hell; if they avoid hell, they will get a barely great life. Z continues to guarantee a very good life alongside a chance of heaven approximately equal to A's chance of avoiding hell. Nevertheless, SD still prefers to chase the small chance of a great life, now risking near-certain hell in its pursuit.[10]

There seems to be an intuitive fix here, whereby avoiding a *horrible* life (at a welfare level of less than or equal to −80) is also assigned superiority over non-great lives. On this version, SD will maximize their probability of living a great life *and* minimize their probability of living a horrible life. Nevertheless, the problem can be reintroduced by assigning A's life in State 1 a welfare level of −80 + ε, which barely avoids the *horrible* range. The problem remains slightly watered down: SD is no longer ignoring *hell*, but merely a very bad life.

Raising the threshold above −80 would dilute the problem further but exacerbate another problem: SD will begin to care too much about disvaluable lives. To illustrate, suppose we raise the threshold all the way to 0 to ensure SD is not insensitive to *any* disvaluable lives. Now, SD will treat avoiding *any* disvaluable life as superior to very good lives; they will refuse to take a one-in-a-million risk of a life that is barely not worth living

---

[10]Parfit appeared to be aware of a related unfortunate result for perfectionism: if enjoying the best things in life takes lexical superiority over all other experiences, then perfectionists will not be very concerned with alleviating suffering. He eventually decided that the matter was extraneous since the Repugnant Conclusion does not involve suffering (Parfit 1986: 163–64).

for the complement probability (near-certainty) of a very good life. So, introducing a range of bad lives that are not inferior to great lives will not escape the problem.

3.6 Accumulation of ignored states

As we saw in 3.2, we can create value that flies under SD's radar by adding a state whose probability is below the threshold. But why stop at one? With each new state in the decision problem, the total probability of the discounted states will agglomerate, but SD will continue to ignore them due to their small individual probabilities. For example, suppose we add one million states to the decision problem, each almost identical to State 3. (To avoid licensing SD to lump all these new states into one, much more probable state,[11] we might assign each of Z's outcomes in these new states a *unique* value approximately equal to 100 and each of A's outcomes a unique value approximately equal to −100. For simplicity, this complication is ignored in Table 6.)

  SD still prefers A because they will ignore the new, insufficiently probable states.[12] Now, A would almost certainly result in hell, while Z would almost certainly result in heaven. For instance, if ε = one in one billion, then the probability of heaven on Z is 0.999, and the probability of hell on A is 0.999999.[13] Both prospects' remaining probabilities are assigned to State 2, in which they would perform approximately equally. (A would produce a life at 80, while Z would produce a life at 80 − ε.)

  We have one more step to get to the Preposterous Conclusion, but we have arrived at what we might call:

> *The Intrapersonal Preposterous Conclusion*: Near-certain hell is better than near-certain heaven (even when both alternative outcomes are approximately equal in value).

3.7 Populations

So far, we have been judging what is better for a single person that might exist. The problems we have identified in sections 3.1–3.6 are magnified when we look at the value of *populations*. We established the Intrapersonal Preposterous Conclusion by

---

[11]While this would bring into focus the problem of partition variance, our task here is to establish the Preposterous Conclusion.

[12]SD can avoid this result by relativizing the threshold to the most probable state (e.g., Lee-Stronach 2018: 801). So, as State 1's probability falls (as it is redistributed), the threshold falls accordingly, and each new state will be above the new threshold. However, we can reintroduce the problem by dividing the probability of each new state by the amount required to stay under the threshold, then multiplying the number of new states by the same number. This group of problematic states will have the same total probability (since the decrease in each state's probability is offset by the addition of the new states), but each individual state will be below the threshold.

[13]It is worth noting that an approach concerned with the probability of a great life on some prospect (rather than the probability of states) would avoid these issues: the probability of Z producing a great life is 0.999, and thus presumably over the threshold (even though the probability of every state *in which* Z produces a great life is minuscule). It is unclear how such an approach would prevent insufficiently probable acquisitions of great lives from contributing to the value of the prospect (since without this feature, agents will chase arbitrarily small probabilities), but perhaps the details could be filled out plausibly. Nevertheless, careful constructions of discounting theories invariably discount states' probabilities. Kosonen does not explicitly make this distinction, but in more precise moments appears to follow suit (e.g., Kosonen 2021: 212, note 35).

Table 6. The Intrapersonal Preposterous Conclusion

	State 1 $(1 - 1/1m - ((1/1m - \varepsilon) \times 1m))$	State 2 $(1/1m)$	State 3 $(1/1m - \varepsilon)$	. . .	State 1,000,002 $(1/1m - \varepsilon)$
A	−100	80	−100	. . .	**−100**
Z	80 − ε	80 − ε	100	. . .	**100**

Table 7. The Preposterous Conclusion

	State 1 $(1 - 1/1m - ((1/1m - \varepsilon) \times 1m))$	State 2 $(1/1m)$	State 3 $(1/1m - \varepsilon)$	. . .	State 1,000,002 $(1/1m - \varepsilon)$
A	$n \times {-100}$	$10b \times 80$	$n \times {-100}$	. . .	$n \times {-100}$
Z	$n \times (80 - \varepsilon)$	$n \times (80 - \varepsilon)$	$n \times 100$	. . .	$n \times 100$

making SD ignore every possible outcome except A's outcome in State 2. Since these outcomes are being ignored due to their lack of great lives or their corresponding states' improbability, increasing the number of people in these outcomes will not affect SD's preference. If it is preposterous to prefer A to Z when there is one person in the equation, it should be more preposterous to maintain this preference when there are more people in the equation. In Table 7, each outcome (other than A's in State 2) contains lives at the same welfare level as in Table 6, but the number of such lives is arbitrarily great (n). We also increase A's outcome in State 2 to ten billion: the number of great lives that is better than any number of mediocre lives. (This step is unnecessary if we are dealing with a *Strong Superiority* theory.)

In preferring A to Z, SD endorses the Preposterous Conclusion:

> *The Preposterous Conclusion*: Near-certain hell for n people (and otherwise a small fraction living barely great lives) is better than near-certain heaven (and otherwise very good lives) for n people.

Just as it is difficult to imagine a worse result for a theory of well-being than the Intrapersonal Preposterous Conclusion, it is difficult to imagine a worse result for a population axiology than the Preposterous Conclusion.

4 Conclusion

Axiologies can avoid the Repugnant Conclusion by endorsing the superiority of great lives over mediocre lives. They can deny Weak Pareto by adopting discounting, allowing them to embrace the Intrapersonal Repugnant Conclusion without thereby accepting the Repugnant Conclusion. But Superiority Discounting avoids the Repugnant Conclusion only at the expense of the Preposterous Conclusion. The Superiority Discounter might patch their theory to reduce the impact of some of these issues, the result of which would be a moderately preposterous conclusion. But giving up all the features that can be exploited would simply be to give up Superiority Discounting.

Competing interests. The author declares none.

References

Aboodi, Ron, Adi Borer, and David Enoch. 2008. Deontology, Individualism, and Uncertainty: A Reply to Jackson and Smith. *The Journal of Philosophy*, **105**: 259–72.

Arrhenius, Gustaf, and Wlodek Rabinowicz. 2015. Value Superiority. In Iwao Hirose and Jonas Olson (eds.), *The Oxford Handbook of Value Theory* (Oxford University Press), pp. 225–48.

Bernoulli, Daniel. 1738. Exposition of a New Theory on the Measurement of Risk. *Econometrica*, **22**: 23–36.

Bjorndahl, Adam, A. J. London, and K. J. S. Zollman. 2017. Kantian Decision Making under Uncertainty: Dignity, Price, and Consistency. *Philosophers' Imprint*, **17**, 1–22.

Borel, Emile. 1962. *Probabilities and Life* (New York: Dover).

Brentano, Franz. 2009. *The Origin of Our Knowledge of Right and Wrong* (London: Taylor & Francis Group).

Buchak, Lara. 2013. *Risk and Rationality* (Oxford: Oxford University Press).

Buffon, Georges-Louis Leclerc de. 1777. Essai D'arithmétique Morale. In *Supplement a l'histoire Naturelle* (Paris: Imprimerie royale).

Chalmers, Adam. 2017. *An Offer You Can't (Rationally) Refuse: Systematically Exploiting Utility-Maximisers with Malicious Gambles*. Honours Thesis. University of Sydney.

Condorcet, Jean-Antoine-Nicolas de Caritat. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix* (Paris: Imprimerie royale).

Crisp, Roger. 1992. Utilitarianism and the Life of Virtue, *The Philosophical Quarterly*, **42**: 139–60.

d'Alembert, Jean Le Rond. 1761. *Opuscules mathématiques, ou Mémoires sur différens sujets de géométrie, de Méchanique, d'optique, d'astronomie &c* (Paris: David).

Edwards, Ren B. 1979. *Pleasures and Pains: A Theory of Qualitative Hedonism* (Ithaca, NY: Cornell University Press).

Glover, Jonathan. 1977. *Causing Death and Saving Lives* (New York: Penguin).

Griffin, James. 1986. *Well-Being: Its Meaning, Measurement, and Moral Importance* (Oxford: Clarendon Press).

Haque, Adil Ahmed. 2012. Killing in the Fog of War. *Southern California Law Review*, **86**: 63–116.

Hawley, Patrick. 2008. Moral Absolutism Defended. *Journal of Philosophy*, **105**: 273–75.

Hutcheson, Francis. 1755. *A System of Moral Philosophy* (London: A. M. Kelley).

Jordan, Jeff. 1994. The St. Petersburg Paradox and Pascal's Wager, *Philosophia*. **23**: 207–22.

Kagan, Shelly. 1989. *The Limits of Morality* (Oxford: Oxford University Press).

Kosonen, Petra. 2021. Discounting Small Probabilities Solves the Intrapersonal Addition Paradox. *Ethics*, **132**: 204–17.

Lazar, Seth. 2017. Risky Killing. *Journal of Moral Philosophy*, **16**: 1–26.

Lee-Stronach, Chad. 2018. Moral Priorities under Risk. *Canadian Journal of Philosophy*, **48**: 793–811.

Lemos, Noah M. 1993. Higher Goods and the Myth of Tithonus. *Journal of Philosophy*, **90**: 482–96.

Mill, J. S. 1998. *Utilitarianism*, ed. Roger Crisp (Oxford: Oxford University Press).

Monton, Bradley. 2019. How to Avoid Maximizing Expected Utility. *Philosophers' Imprint*, **19**: 1–25.

Nebel, Jacob M. 2019. An Intrapersonal Addition Paradox. *Ethics*, **129**: 309–43.

Parfit, Derek. 1986. Overpopulation and the Quality of Life. In Peter Singer (ed.), *Applied Ethics* (Oxford: Oxford University Press), pp. 145–64.

Robert, David. 2018. Expected Comparative Utility Theory: A New Theory of Rational Choice, *Philosophical Forum*, **49**: 19–37.

Ross, W. D. 1930. *The Right and the Good: Some Problems in Ethics* (Oxford: Clarendon Press).

Schwitzgebel, Eric. 2017. 1% Skepticism. *Noûs*, **51**: 271–90.

Skorupski, John. 1999. *Ethical Explorations* (Oxford University Press).

Smith, Nicholas J. J. 2014. Is Evaluative Compositionality a Requirement of Rationality? *Mind*, **123**: 457–502.

Smith, Nicholas J. J. 2016. Infinite Decisions and Rationally Negligible Probabilities. *Mind*, **125**: 1199–212.

Tarsney, Christian. 2018. Moral Uncertainty for Deontologists. *Ethical Theory and Moral Practice*, **21**: 505–20.