

Right Research

Modelling Sustainable Research Practices
in the Anthropocene

*Edited by Chelsea Miya, Oliver Rossier
and Geoffrey Rockwell*



<https://www.openbookpublishers.com>

© 2021 Chelsea Miya, Oliver Rossier and Geoffrey Rockwell. Copyright of individual chapters is maintained by the chapter's author.



This work is licensed under a Creative Commons Attribution 4.0 International license (CC BY 4.0). This license allows you to share, copy, distribute and transmit the work; to adapt the work and to make commercial use of the work providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work). Attribution should include the following information:

Chelsea Miya, Oliver Rossier and Geoffrey Rockwell (eds), *Right Research: Modelling Sustainable Research Practices in the Anthropocene*. Cambridge, UK: Open Book Publishers, 2021. <https://doi.org/10.11647/OBP.0213>

Copyright and permissions for the reuse of many of the images included in this publication differ from the above. This information is provided in the captions and in the list of illustrations.

In order to access detailed and updated information on the license, please visit <https://doi.org/10.11647/OBP.0213#copyright>. Further details about CC BY licenses are available at <http://creativecommons.org/licenses/by/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

Digital material and resources associated with this volume are available at <https://doi.org/10.11647/OBP.0213#resources>

Every effort has been made to identify and contact copyright holders and any omission or error will be corrected if notification is made to the publisher.

ISBN Paperback: 9781783749614

ISBN Hardback: 9781783749621

ISBN Digital (PDF): 9781783749638

ISBN Digital ebook (epub): 9781783749645

ISBN Digital ebook (mobi): 9781783749652

ISBN Digital (XML): 9781783749669

DOI: 10.11647/OBP.0213

Cover image by Leanne Olson, *The Clay* at Ryley, CC-BY-NC-ND.

Cover design by Emilie St-Hilaire.

1. Why Should We Try to Be Sustainable?

Expected Consequences and the Ethics of Making an Indeterminate Difference

Howard Nye

Why should we refrain from doing things that, taken collectively, are environmentally destructive, if our individual acts seem almost certain to make no difference? According to the expected consequences approach, we should refrain from doing these things because our individual acts have small risks of causing great harm, which outweigh the expected benefits of performing them. Several authors have argued convincingly that this provides a plausible account of our moral reasons to do things like vote for policies that will reduce our countries' greenhouse gas emissions, adopt plant-based diets, and otherwise reduce our individual emissions. But this approach has recently been challenged by authors like Bernward Gesang and Julia Nefsky. Gesang contends that it may be genuinely impossible for our individual emissions to make a morally relevant difference. Nefsky argues more generally that the expected consequences approach cannot adequately explain our reasons not to do things if there is no precise fact of the matter about whether their outcomes are harmful. In the following chapter, author Howard Nye defends the expected consequences approach against these objections. Nye contends that Gesang has shown at most that our emissions could have metaphysically indeterministic effects

that lack precise objective chances. He argues, moreover, that the expected consequences approach can draw upon existing extensions to cases of indeterminism and imprecise probabilities to deliver the result that we have the same moral reasons to reduce our emissions in Gesang's scenario as in deterministic scenarios. Nye also shows how the expected consequences approach can draw upon these extensions to handle Nefsky's concern about the absence of precise facts concerning whether the outcomes of certain acts are harmful. The author concludes that the expected consequences approach provides a fully adequate account of our moral reasons to take both political and personal action to reduce our ecological footprints.

1. Environmental Destruction and the Ethics of Collective Action

Why should we try to reduce our destructive impacts on the environment, when it can seem that the effects of our individual acts are too small to make an ethically important difference? As Walter Sinnott-Armstrong puts the challenge, why, for instance, should one seek to reduce one's emissions of greenhouse gases [GHGs] by cycling and taking public transit instead of driving, if it seems that 'Climate change occur[s] on such a massive scale that my individual driving makes no difference to the welfare of anyone'?¹

This is an instance of a general ethical problem about collective action, which is of great practical as well as theoretical importance. The view that our own reductions of GHG emissions will have too small of an effect to make an important difference appears to be the last line of defense of those inclined to oppose action to address climate change—if at any point they do tire of denying the overwhelming

1 Walter Sinnott-Armstrong, 'It's not my fault: Global warming and individual moral obligations', in *Perspectives on Climate Change: Science, Economics, Politics, Ethics*, ed. by Walter Sinnott-Armstrong and Richard Howarth (Amsterdam: Elsevier, 2005), pp. 285–307 (p. 301), [https://doi.org/10.1016/S1569-3740\(05\)05013-3](https://doi.org/10.1016/S1569-3740(05)05013-3).

evidence that the climate is changing, that the change is anthropogenic, and that the change is extremely harmful.² But even those who accept the overwhelming evidence for harmful anthropogenic climate change and agree that we should do something about it at the level of social policy can (like Sinnott-Armstrong) be sorely tempted—including as a rationalization for personal inaction—by the thought that individual attempts to act in less environmentally destructive ways are futile.

The theoretical importance of such problems of collective action concerns whether in general there are moral reasons in these cases for individuals to act, and what sort of moral theory best accounts for this. For instance, according to the

Expected consequences approach: in *collective action cases* where our acts together are collectively harmful (or beneficial), our individual acts do in fact have a chance [often small] of causing harm (or benefit) [often large] to others. As such, the moral importance of avoiding this risk of harm (or securing this chance of benefit) typically outweighs the possible benefits to us of performing (or failing to perform) these acts.³

-
- 2 For an authoritative guide to this overwhelming evidence see IPCC 2014, *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. by Core Writing Team, R.K. Pachauri and L.A. Meyer (Geneva: IPCC, 2014), https://www.ipcc.ch/site/assets/uploads/2018/05/SYR_AR5_FINAL_full_wcover.pdf. For discussion of the ‘it won’t make a difference if we reduce our emissions’ objection by those who oppose action on climate change, see e.g. Michael Mann, *The Hockey Stick and the Climate Wars: Dispatches from the Front Lines* (New York: Columbia University Press, 2012), <https://doi.org/10.7312/mann15254>. A vivid recent example of this rationale for inaction is the National Highway Traffic Safety Administration’s defense of its fuel efficiency rollbacks on the grounds that other actors’ GHG emissions will cause harmful climate change regardless of what the NHTSA does (National Highway Traffic Safety Administration, *Draft Environmental Impact Statement for The Safer Affordable Fuel-Efficient (SAFE) Vehicles Rule for Model Year 2021–2026 Passenger Cars and Light Trucks* (2018), pp. 5–30, https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/ld_cafe_my2021-26_deis_0.pdf).
- 3 Defenders of this approach include Peter Singer, ‘Utilitarianism and vegetarianism’, *Philosophy and Public Affairs*, 9 (1980), 325–337 (pp. 335–336); Derek Parfit, *Reasons and Persons* (New York: Oxford University Press, 1984), pp. 73–86, <https://doi.org/10.1093/019824908x.001.0001>; Gaverick Matheny, ‘Expected utility, contributory causation, and vegetarianism’, *Journal of Applied Philosophy*, 19 (2002), 293–297, <https://doi.org/10.1111/1468-5930.00223>; Alastair Norcross, ‘Puppies, pigs, and people: Eating meat and marginal cases’, *Philosophical Perspectives*, 18 (2004), 229–245 (pp. 231–233), <https://doi.org/10.1111/j.1520-8583.2004.00027.x>; Shelly Kagan,

If this is correct, then we can explain our moral reasons to omit being complicit in harmful practices and contribute to beneficial practices in terms of familiar principles of non-maleficence not to harm others, beneficence to benefit others, and responsible decision-making under conditions of uncertainty. We need not add any fundamentally distinct ethical principles to those acknowledged by moral theories like Rossian pluralism, according to which there is a plurality of basic moral reasons to act, including non-maleficence and beneficence;⁴ and act consequentialism, according to which our moral reasons to do something are proportional simply to the amount of good it will do.⁵ Nor need we find any fundamentally new principles of responsible decision-making under conditions of uncertainty beyond the teachings of standard decision theory, according to which (very roughly) we should weigh in favour of an act the benefits it may bring in proportion to their magnitude and likelihood of occurring (its ‘expected benefits’), and weigh against it the harms it may produce in proportion to their magnitude and likelihood of occurring (its ‘expected harms’).⁶

Advocates of the expected consequences approach have argued convincingly that it provides a plausible account of the moral case

‘Do I make a difference?’, *Philosophy and Public Affairs*, 39 (2011), 105–141, <https://doi.org/10.1111/j.1088-4963.2011.01203.x>; Avram Hiller, ‘Climate change and individual responsibility’, *The Monist*, 94 (2011), 349–368, <https://doi.org/10.5840/monist201194318>, and ‘Morally significant effects of ordinary individual actions’, *Ethics, Policy and Environment*, 14 (2011), 19–21, <https://doi.org/10.1080/21550085.2011.561588>; John Broome, *Climate Matters: Ethics in a Warming World* (New York: W. W. Norton & Co., Inc., 2012), pp. 73–78; and Christopher Morgan-Knapp and Charles Goodman, ‘Consequentialism, climate harm, and individual obligations’, *Ethical Theory and Moral Practice*, 18 (2015), 177–190, <https://doi.org/10.1007/s10677-014-9517-9>.

- 4 See e.g. W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930), pp. 16–64.
- 5 See e.g. Henry Sidgwick, *The Methods of Ethics*, 7th edn (London: Macmillan and Co., Limited, 1907), pp. 337–361 (373–390, and 418–459); and Shelly Kagan, *The Limits of Morality* (Oxford: Oxford University Press, 1989), <https://doi.org/10.1093/0198239165.001.0001>.
- 6 Classic statements of which include Frank P. Ramsey, ‘Truth and probability’ (1926), reprinted in *Decision, Probability, and Utility: Selected Readings*, ed. by Peter Gärdenfors and Nils-Eric Sahlin (New York: Cambridge University Press, 1988), pp. 19–47, <https://doi.org/10.1017/cbo9780511609220>; John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton: Princeton University Press, 1944), pp. 15–31, <https://doi.org/10.1515/9781400829460>; and Leonard J. Savage, *The Foundations of Statistics* (New York: Dover Publications, Inc., 1954).

for individual acts that can help to reduce our ecological footprints. As I discuss below, these include voting for policies that will reduce our countries' GHG emissions, eating plant-based diets, reducing our air-travel and car-travel by substituting e-conferences for in-person conferences, and walking, cycling, and taking public transit instead of driving. But this account has recently been challenged by Bernward Gesang, who contends that it may be not only unlikely but genuinely impossible for our individual GHG emissions to make a difference to morally important outcomes.⁷ Moreover, Julia Nefsky has challenged the general adequacy of the expected consequences approach by arguing in effect that it cannot explain our reasons not to do things in certain collective action cases where there is no precise fact of the matter about whether their outcomes are harmful.⁸

In this chapter I defend the expected consequences approach against these challenges. I argue that Gesang has shown at most that our emissions could have metaphysically indeterministic effects that lack precise objective chances. But the expected consequences approach has been extended to cases of indeterminism and imprecise probabilities by authors like Krister Bykvist and Susanna Rinard.⁹ I show how these extensions vindicate the application of the expected consequences approach to the scenario that Gesang attempts to describe. Moreover, I argue that these extensions of the expected consequences approach can be used to respond to Nefsky's challenge by appropriately explaining our reasons not to do things in collective action cases where there are no precise facts about whether their outcomes are harmful. I conclude that these extensions of the expected consequences approach enable it to provide a fully adequate account of our moral reasons to take both political and personal action to reduce our ecological footprints.

7 Bernward Gesang, 'Climate change—do I make a difference?', *Environmental Ethics*, 39 (2017), 3–19, <https://doi.org/10.5840/enviroethics20179261>.

8 Julia Nefsky, 'Consequentialism and the problem of collective harm: A reply to Kagan', *Philosophy and Public Affairs*, 39 (2012), 364–395, <https://doi.org/10.1111/j.1088-4963.2012.01209.x>.

9 Krister Bykvist, 'Normative supervenience and consequentialism', *Utilitas*, 15 (2003), 27–49, <https://doi.org/10.1017/s0953820800003757>; Susanna Rinard, 'A decision theory for imprecise probabilities', *Philosophers' Imprint*, 15 (2015), 1–16, <https://hdl.handle.net/2027/spo.3521354.0015.007>.

2. The Expected Consequences Approach to Collectively Destructive Acts

One of the great advantages of the expected consequences approach is that it explains our moral reasons in collective action cases in terms of extremely plausible general principles of non-maleficence, beneficence, and standard decision theory's injunction to weigh acts' expected harms against their expected benefits. These entail that there can be a decisive moral case against performing an act because it carries a small risk of causing a great deal of harm in return for relatively small expected benefits. This seems to be an excellent explanation of why there is a decisive moral case against such acts as speeding through residential areas when late for work, bouncing a ball around a nuclear missile's launch button just for fun, and shooting into an occupied building just for target practice.

At the same time, these principles provide an extremely plausible and helpful way to determine when an act that would otherwise be wrong in virtue of its risks becomes permissible, namely when its expected benefits are great enough in comparison to its expected harms.¹⁰ There is, for instance, no decisive moral case against—but plausibly one in favour of—such acts as driving at high speed through a residential area and shooting into an occupied building if these are the only ways to rush someone in need of urgent care to the hospital or diffuse a bomb about to explode and kill someone.

The expected consequences approach can be applied straightforwardly to the collective action case of voting. Consider the *Simpsons*-inspired case of

Voting for Kang. It is revealed that one of the major party candidates for president in the United States is actually an evil alien named

10 This can be true whether or not one subscribes to the view that there are constraints on harming, or that certain harmful upshots of our conduct are in themselves harder to justify than certain of their failures to have beneficial upshots. One must only avoid the extremely implausible view that there is an absolute prohibition on harmful upshots that translates into an absolute injunction to avoid any risk of harm no matter how small in return for any chance of benefit no matter how great (see e.g. See Kagan (1989), pp. 87–91; and Frank Jackson and Michael Smith, 'Absolutist moral theories and uncertainty', *The Journal of Philosophy*, 103 (2006), 267–283, <https://doi.org/10.5840/jphil2006103614>).

Kang who will enslave and torture everyone in the country (with no benefits to anyone other than US citizens) if and only if he wins the election. The other candidate (unlike in ‘Citizen Kang’) is a typical politician who will maintain the status quo. Kang has, however, rigged voting booths to pay out \$50 to everyone who votes for him.

It seems clear that one has decisive moral reason to forgo the \$50 and omit voting for Kang—and indeed vote against him. As Derek Parfit showed, this intuition can be justified by the expected consequences approach.¹¹ Updating Parfit’s calculations, there are about 326 million US citizens, and the average probability of one’s vote deciding a presidential election is one in 60 million.¹² The expected harm to others of voting for Kang is thus $(1/60 \text{ million}) \times h \times (326 \text{ million}) = 5.43 \times h$, where h is the per-citizen harm inflicted by Kang. Very clearly, the certain benefit of \$50 to you (which we may assume can be spent before Kang takes power if he wins) is absolutely trivial in comparison to the per-individual harm inflicted by Kang on one other individual—let alone an expected harm more than five times as great.

To appreciate the expected benefit of voting for sustainable policies, we can consider a more realistic scenario of an election between Superior and Inferior. If Superior wins, she will implement a Green New Deal that will reduce the emissions of the US to net zero by 2050 while benefitting most workers by providing jobs and stimulating aggregate demand. If Inferior wins she will maintain the status quo on US emissions. While the Green New Deal’s economic benefits would

11 Parfit (1984), pp. 73–74.

12 Andrew Gelman, Nate Silver and Aaron Edlin, ‘What is the probability that your vote will make a difference?’, *Economic Inquiry*, 50 (2012), 321–326, <https://doi.org/10.1111/j.1465-7295.2010.00272.x>; see also Pierre-Antoine Kremp and Andrew Gelman, ‘What is the chance that your vote will decide the election?’, https://pkremp.github.io/pr_decisive_vote.html. Depending upon one’s state of residence, one’s vote will have a greater or lesser chance of deciding the election. For instance, if one votes in Colorado, Michigan, Connecticut or Wyoming, one’s chance of deciding the election will be respectively one in 1 million, 3 million, 40 million or 30 billion. This, of course, assumes that voting patterns in the election between Kang and his opponent would follow those of other elections. One could, purely hypothetically, imagine this being so due to a distribution of propensities to be more concerned about getting the \$50 or more concerned about voting against Kang that are isomorphic to current partisan voting patterns.

be substantial, expected differences in lives lost due to climate policy alone are likely to be enormous. For instance, as I will discuss more below, John Nolt estimates that the expected harm of an average US citizen's current lifetime GHG emissions is 1–2 human lives lost.¹³ Since Superior will zero-out the emissions of all 326 million US citizens, the expected benefit of voting for her due simply to her climate policy will be one's 1/60 million chance of deciding the election times the 326 to 652 million lives that can be expected to be saved by this policy, or 5.43 to 10.87 lives. Very clearly, the costs to oneself of voting are absolutely trivial in comparison to the moral importance of saving more than five to ten lives!

Indeed, even if one lives in Wyoming and we take the more conservative estimate of Superior's saving only 326 million lives, the expected benefit of one's voting for her would still be $(1/30 \text{ billion}) \times (326 \text{ million}) = 0.0109$ lives which (assuming a life expectancy of 80 years) is 317.31 days of life. Again, the cost to oneself of voting seems completely trivial in comparison to someone else's being deprived of 317.31 days of life. Even if Superior would, rather than eliminate US emissions, simply reduce them by 10% or just 1%, the expected benefits of a vote for Superior in Wyoming due simply to this policy would still be 31.73 or 3.17 days of life. Since the costs to oneself of voting are clearly less important than someone's not being deprived of a month or even three days of life, it would still clearly be worthwhile to vote for Superior, even in Wyoming.¹⁴

13 John Nolt, 'How harmful are the average American's greenhouse gas emissions?', *Ethics, Policy and Environment*, 14 (2011), 3–10, <https://doi.org/10.1080/21550085.2011.561584>.

14 Those who are inclined to find it too fanciful for an election to be decided by a single vote should consider the very real-world case of the 2017 election in Virginia's 94th House of Delegates district, where the vote was exactly tied, and control of the House depended upon a single seat. See e.g. Emily Tillet, 'Virginia election results 2017: Republican David Yancey wins Virginia House seat', *CBS News* (January 4, 2018), <https://www.cbsnews.com/news/virginia-election-results-lottery-drawing-house-of-delegates-david-yancy-winner-virginia-house-seat/>. Had one more voter voted for the Democrat, the seat would have been won, and instead of the Republicans winning a drawing by lot that led to them holding a 51 to 49 seat majority, control of the House would have been shared. There are records of at least 64 similar elections between 1822 and 2018, 21 of which were tied and would have been decided by a single additional vote, and 43 of which were in fact decided by a single vote. See references to records compiled at 'List of close election results',

Voting is a paradigm example of what Shelly Kagan calls a *triggering case*, where if a certain threshold of individual contributions is crossed this triggers a morally important effect. In such cases, if one's act is part of a cohort that falls short of the triggering number or adds to the surplus above the triggering number, then it (as well as the other members of the cohort) makes no difference to the effect. But if one's act is part of a cohort that exactly crosses the threshold and triggers the effect, then it (as well as the other acts) makes all of the difference to it. To a first approximation voting is a triggering case where one's act has a chance of crossing only a single triggering threshold.¹⁵ There are, however, many important cases in which one's act has a chance of being a part of several cohorts each of which exactly cross different thresholds that trigger morally important effects.

For instance, as Peter Singer, Gaverick Matheny, Alastair Norcross, and Kagan argue, purchasing animal products instead of plant-based alternatives is an instance of a triggering case where one's purchase has a small chance of crossing many thresholds, each of which would result in a great deal of harm to animals and the environment.¹⁶ Given the price mechanism in a market economy, there must be some number N (e.g. 10,000) of additional chicken purchases that causes N more chickens to be bred, tortured, and killed—at the expense of much more land, grain, and water inputs, polluting waste, and unsequestered GHG

Wikipedia, The Free Encyclopedia (October 9, 2020), https://en.wikipedia.org/wiki/List_of_close_election_results.

15 In reality, votes can matter beyond their effect on actually electing candidates by adding to or taking away from margins of victory or defeat, which can affect the extent to which politicians take themselves to have mandates, take certain policies to be favoured by likely voters, and take themselves to be likely to be elected or re-elected if they behave in the way the winning or losing candidate did. How individual votes affect perceptions of margins might not be straightforward, but one way this could work is for votes to have chances of creating totals that round to certain significant digits that stick in the minds of the relevant decision makers in ways they would not have had they been one vote short of this rounding threshold. Since there are presumably many different such totals that matter to many different such decisions makers, voting is to a more accurate approximation a triggering case with one triggering threshold of great significance and several other triggering thresholds of somewhat lesser significance.

16 Singer (1980), pp. 335–336; Matheny (2002); Norcross (2004), pp. 232–234; and Kagan (2011), pp. 110–127.

emissions—than the consumption of plant-based alternatives.¹⁷ In *N-1*

-
- 17 On the torture of chickens in the meat industry and the much greater inputs and polluting outputs of animal agriculture, see e.g. Peter Singer and Jim Mason, *The Ethics of What We Eat: Why Our Food Choices Matter* (New York: Rodale Press, 2006); Richard Oppenlander, *Food Choice and Sustainability: Why Buying Local, Eating Less Meat, and Taking Baby Steps Won't Work* (Minneapolis: Langdon Street Press, 2013); Food and Agriculture Organization of the United Nations, *Tackling Climate Change Through Livestock: A Global Assessment of Emissions and Mitigation Opportunities* (Rome: FAO, 2013), <http://www.fao.org/3/i3437e/i3437e00.htm>; Joseph Poore and Thomas Nemcek, 'Reducing food's environmental impacts through producers and consumers', *Science* 360 (2018), 987–992, <https://doi.org/10.1126/science.aaq0216>; and Alon Shepon et al., 'The Opportunity Cost of Animal Based Diets Exceeds All Food Losses', *Proceedings of the National Academy of Science [PNAS]* 115(15) (2018), 3804–3809, <https://doi.org/10.1073/pnas.1713820115>.

If one is tempted to think that bringing chickens into existence and killing them is justified on the grounds that it benefits the chickens on net, one should note that: (i) due to the much greater ecological use and damage of producing chickens instead of plant-based alternatives, far more wild animals who would likely live much better lives are caused to exist by the consumption of plant-based alternatives, and (ii) the view that it is OK to support the practice of bringing beings into existence and killing them on the grounds that this benefits them on net entails that one would be justified in supporting similar practices of bringing into existence human children for the purpose of torturing or killing them on the grounds that this would benefit them on net. More plausibly, coming into existence is not a morally important benefit but a precondition for morally important benefit and harm, and bringing someone into this state gives one a special obligation to ensure that she will not come to harm such that if one cannot discharge this obligation one should not bring her into existence. For related ideas about the ethics of causing beings to exist and the genesis of special obligations, see Melinda Roberts, 'The asymmetry: A solution', *Theoria*, 77 (2011), 333–367, <https://doi.org/10.1111/j.1755-2567.2011.01117.x>; and Seana Shiffrin, 'Wrongful life, proactive responsibility, and the significance of harm', *Legal Theory*, 5 (1999), 117–148. <https://doi.org/10.1017/s1352325299052015>.

Producers make production decisions on the basis of the anticipated future prices of products (as well of course as costs of production including the opportunity costs of inputs), a large input to which are present prices. One way in which one's purchases can affect the price of a product and thus the likelihood of greater amounts of it being produced is explained by the theory of marginal pairs pioneered by Eugen Böhm-Bawerk, *The Positive Theory of Capital*, trans. by William Smart (1889, reprinted in New York: G. E. Stechert & Co., 1930), pp. 198–213; and developed by Murray Rothbard, *Man, Economy, and the State* (Princeton Van Nostrand, 1962; reprinted in Auburn: Ludwig von Mises Institute, 2001), pp. 106–126. According to this theory a product's equilibrium price is determined to be somewhere in the range between, on the high end: (1) the lesser value of (a) the greatest amount that the *last buyer* (who is willing to pay the least for the product among those who actually buy it) is willing to pay for it, and (b) the lowest amount that the *first excluded seller* (who is willing to accept the least for the product among those who are unwilling to sell it) would have been willing to accept for it; and, on the low end, (2) the greater of (c) the lowest amount that the *last seller* (who is least willing to accept less for the product among those who actually sell it) would have

out of N cases, one's purchase will not make the crucial difference between being part of a cohort that does not as opposed to does cause another N chickens to be destructively tortured and killed, but in 1 out of N cases one's act will make this crucial difference. So the expected harm of one's act of purchasing a chicken is $(1/N) \times N \times h = h$, where h is the harm done to a chicken by being tortured and killed together with the environmental cost of doing this to her. The act thus has the same expected harm, and is as morally important to avoid, as directly torturing and killing one chicken and causing the attendant ecological damage for the mere taste pleasure of eating her.¹⁸

Animal products are responsible for at least 14.5% of GHG emissions—more than all transportation exhaust—and eliminating them from our diets is the easiest and most effective single thing that most of us can do to reduce our carbon footprints.¹⁹ John Broome, Avram

been willing to accept for it, and (d) the greatest amount that the *first excluded buyer* (who is willing to pay the most for the product among those who are unwilling to buy it) would have been willing to pay for it. If one would have been the last buyer had one bought the product, then one's abstention may (i) decrease the upper limit of the equilibrium price by causing the former first excluded buyer to become the new last buyer who is willing to pay the least for the product, and thus also (ii) decrease the lower limit of the equilibrium price by causing the former second excluded buyer to become the new first excluded buyer who is only willing to pay less for the product.

- 18 The expected benefits are at most taste pleasure, since plant-based diets are at least as healthy (and in practice often healthier—see e.g. Academy of Nutrition and Dietetics, 'Position of the Academy of Nutrition and Dietetics: Vegetarian diets', *Journal of the Academy of Nutrition and Dietetics*, 116 (2016), 1970–1980, <https://doi.org/10.1016/j.jand.2016.09.025>)—and equally nutritious plant-based substitutes like legumes and grains are actually less expensive (cf. e.g. Mary Flynn and Andrew Schiff, 'Economical healthy diets (2012) including lean animal protein costs more than using extra virgin olive oil', *Journal of Hunger & Environmental Nutrition*, 10 (2015), 467–482, <https://doi.org/10.1080/19320248.2015.1045675>).
- 19 See e.g. FAO UN, *Tackling Climate Change Through Livestock*; and Gidon Eshel and Pamela Martin, 'Diet, energy, and global warming', *Earth Interactions*, 10 (2006), 1–17, <http://doi.org/10.1175/ei167.1>. Eshel and Martin find that on average going vegan reduces one's carbon footprint by about 1.5 tonnes CO₂ equivalent, about 50% more than switching from a standard car to a hybrid, which reduces it on average by about 1 tonne. 14.5% is actually an extremely conservative estimate of the contribution of animal agriculture to anthropogenic climate change, since among other things it uses a low estimate of the effects of methane emissions from livestock and underestimates the destructive impacts of animal agriculture on carbon sinks; see e.g. Julie Wolf, Ghassem Asrar and Tristram West, 'Revised methane emissions factors and spatially distributed annual carbon fluxes for global livestock', *Carbon Balance and Management*, 12 (2017), 1–24, <https://doi.org/10.1186/s13021-017-0084-y>; Robert Goodland and Jeff Anhang, 'Livestock and climate

Hiller, Christopher Morgan-Knapp, and Charles Goodman have argued convincingly that other instances of emitting GHGs are also triggering cases with many chances of triggering harmful outcomes. As Broome observes:

Greenhouse gas harms people in multifarious ways. Each of them is chancy to some extent. A particular storm will be harmful only if the water rises above the flood defenses. Each increase in the amount of greenhouse gas in the air slightly increases the quantity of rain, but it will be a matter of chance whether the particular quantity of gas you emit this year will be enough to cause a flood on any particular occasion. Your emission increases the likelihood of a flood, but it might not actually cause any particular flood... But during the centuries they are in the air they will have the chance of causing harm on innumerable occasions. It is extraordinarily unlikely that they will do no harm at all.²⁰

Broome is discussing our lifetime emissions rather than, say, the emissions from a particular drive. But even if the emissions from a particular drive are likely to do no harm at all, they do, as Morgan-Knapp and Goodman argue, have a small chance of causing a huge amount of damage through dramatic ‘butterfly effects’ that cascade into harmful events like storms, floods, droughts, and heat-waves. Moreover, because climate science dictates that acts which emit GHGs are more likely to have these butterfly-effect-caused harms and no more likely to have butterfly-effect-caused benefits than their omission, these small chances of great harm are not counterbalanced by equally sized butterfly-effect-caused benefits and do not ‘drop out’ of the decision theoretic evaluation.²¹

change: What if the key actors in climate change are... cows, pigs and chickens?’, *WorldWatch*, November/December 2009, <https://awellfedworld.org/wp-content/uploads/Livestock-Climate-Change-Anhang-Goodland.pdf>; and Timothy Searchinger et al., ‘Assessing the Efficiency of Changes in Land Use for Mitigating Climate Change’, *Nature* 564 (2018), 249–253, <https://doi.org/10.1038/s41586-018-0757-z>.

20 Broome (2012), p. 76.

21 Unlike the completely random and non-directional risks of butterfly-effect-caused harms from any act, which are counterbalanced by equal chances of butterfly-effect-caused benefits and can thus be ignored—at least so long as there are no constraints on harming (see Howard Nye, ‘Chaos and Constraints’, in *Dimensions of Moral Agency*, ed. by David Boersema (Newcastle: Cambridge Scholars Publishing 2014), <https://www.cambridgescholars.com/resources/pdfs/978-1-4438-6692-7-sample.pdf>; Morgan-Knapp and Goodman (2015), pp. 183–286).

As Hiller observes, because an emission's chance of crossing thresholds and triggering harmful effects is proportional to its size, and we know of nothing else that makes any given emission more or less likely to do this, we can determine an emission's expected harm by multiplying the total expected harm of our collective emissions by the ratio of the emission's size to that of the collective amount.²² John Nolt employs this method to estimate the expected harm of the 1,840 tonnes of CO₂ equivalent emitted by an average US citizen over the course of her life to be the death and/or severe suffering of 1–2 humans.²³ Using the more conservative estimate of one life lost per 1,840 tonnes CO₂ equivalent, Hiller estimates the expected harm of a 40 km Sunday joyride in a car, which emits 14.1kg of CO₂ equivalent, as that of ruining someone's afternoon (or depriving her of 5.37 hours of life). While weighty enough to decisively outweigh the benefits of joyrides, and to strongly favour walking, cycling, and taking public transit over commuting by car, these expected harms can plausibly be outweighed by such sufficiently serious considerations as the need to rush someone to the hospital.

A fuller account of the expected harms of our GHG emissions would take into account their effects on non-human animals. Brian Tomasik argues that, on a conservative estimate, there are at least about 14 wild land vertebrates and 1,400 wild marine vertebrates for every human.²⁴ These individuals are clearly sentient,²⁵ and likely to be at least as vulnerable to the harms of climate change as humans. So a fuller but still conservative estimate of the expected harms of the emissions from a Sunday joyride might also include a proportionally great harm to 14 land vertebrates and 1,400 marine vertebrates. The figure of depriving the human of 5.37 hours of life is based on a full human life expectancy

22 Hiller (2011), 357–358.

23 Nolt (2011).

24 Brian Tomasik, 'How many wild animals are there?', 2018, <http://reducing-suffering.org/how-many-wild-animals-are-there/>.

25 See e.g. David Edelman and Anil Seth, 'Animal consciousness: A synthetic approach', *Trends in Neurosciences*, 32 (2009), 476–484, <https://doi.org/10.1016/j.tins.2009.05.008>; Jaak Panksepp, *Affective Neuroscience: The Foundations of Human and Animal Emotions* (Oxford: Oxford University Press, 1998), Jonathan Balcombe, *Pleasurable Kingdom: Animals and the Nature of Feeling Good* (London: Macmillan, 2006); and Victoria Braithwaite, *Do Fish Feel Pain?* (Oxford: Oxford University Press, 2010), <https://doi.org/10.1086/656881>.

of 80 years. So if, for instance, each of the wild animals had a life expectancy absent climate harm of three years, a proportional expected harm to each of them from the emission of 14.1 kg of CO₂ equivalent might be the deprivation of $3/80 \times 5.37 \times 60 = 12.08$ minutes of life.

Hiller's method can also be used to estimate the moral importance of reducing our air travel, for instance by substituting online video econferences, e-colloquia, and e-symposia for traditional academic gatherings that require flying. It has been estimated that up to 1/3 of the GHG emissions of an institution like the University of California, Santa Barbara are due to air travel, and that a single round-trip continental flight can emit a full tonne of CO₂ equivalent, consuming one's entire carbon budget needed to limit global temperature rise to 1.5 degrees C.²⁶ The University of Alberta's Kule Institute for Advanced Study estimates that, for each year that it has held its Around the World econference in the place of a comparable traditional conference, it has eliminated the need for 200 flights and prevented the emissions of 300 tonnes of CO₂ equivalent.²⁷ The expected benefit of doing this is 300/1,840 of the benefit of saving the full lives of a human, 14 land vertebrates, and 1,400 marine vertebrates, or omitting to end the lives of an elderly human with a life expectancy of about 13 years and 14 land together with 1,400 marine vertebrates with life expectancies of about 6 months. According to the expected consequences approach, the moral reasons to transition to an econference are thus comparable in strength to those in favor of omitting to end the lives of a 69-year-old Canadian, the 14 elderly dogs for whom he cares at his animal sanctuary, and the 1,400 elderly fish living in the lake next to the sanctuary. Since it would be morally imperative not to kill these individuals just to enable us to have a conference, it is at least as clearly imperative for us to bear the much lesser cost—which may actually be a net benefit when we account for the savings of money and time, and the greater accessibility to scholars of lesser means around the world—of substituting an econference like Around the World for a comparable in-person conference.

26 Ken Hiltner, 'A nearly carbon-neutral conference model', <https://hiltner.english.ucsb.edu/index.php/ncnc-guide/>.

27 Trevor Chow-Fraser, Chelsea Miya and Oliver Rossier, *Moving Ideas without Moving People: How to Econference at the University of Alberta* (2018), p. 7, <https://aroundtheworld.ualberta.ca/wp-content/uploads/2018/05/E-Conferencing-Toolkit.pdf>.

3. Gesang's and Nefsky's Challenges: Problems of Indeterminacy

Although he is sympathetic to a similar account of the moral case for reducing our emissions, Gesang has recently argued that the expected consequences approach's account as described above does not straightforwardly succeed. His main contention is that there is a reasonably likely climate scenario in which it is not only unlikely but genuinely impossible for small emissions to cause morally significant effects.²⁸ To make this plausible, Gesang draws an analogy to the sorites paradox for vague properties like that of being a heap, which is that two very plausible views about the property lead to the very implausible conclusion that nothing has the property:

- *Some non-heaps*: a collection of 1 grain of sand is not a heap.
- *Tolerance thesis*: for any number n of grains of sand, if a collection of n grains is not a heap, then a collection of $n+1$ grains is not a heap.
- *Therefore, No heaps*: all collections of grains of sand, no matter how many, fail to be heaps.

28 Gesang (2017), pp. 14–19. Gesang sketches another scenario in which the expected consequences approach would not (at least in the straightforward way suggested by Hiller, Morgan-Knapp, and Goodman) support reducing our emissions, namely one in which there is exactly one triggering threshold and we know with certainty (or at least sufficient probability that the expected harm of our emissions is less than the expected harm of the large amount of emissions crossing the threshold times the ratio of our emissions to the large amount of emissions) that the threshold has been crossed and the continued existence of mammals will be impossible, regardless of whether or not we emit any given quantity of GHGs (pp. 8–9). While Gesang thinks this scenario is very unlikely and does not rest much weight on his discussion of it, it is worth noting that (i) even if we knew with certainty that a threshold had been crossed that would soon make life for mammals impossible, it does not follow that our emissions can have no further morally relevant effects (e.g. for the mammals and others before life becomes impossible, and for the sentient non-mammals after life becomes impossible for the mammals), and (ii) even if we thought it was extremely likely that the threshold would be crossed regardless of whether we emit, the tiniest probability of our emissions influencing this catastrophic event could easily, in virtue of the enormous size of the harm, be enough to make our emissions' expected harms outweigh their expected benefits. Compare this to the case of voting for president in Wyoming discussed above.

As Gesang notes, most theorists view vagueness and the consequent Sorites paradox as resulting from the semantic or mental feature that there is no precise fact of the matter about what falls under the concepts signified by vague predicates like 'is a heap' and 'is bald'. His problem is thus to explain how there could be a metaphysical sorites paradox in the case of GHG emissions, in which it seems plausible that small additional GHG emissions can never make the difference between the absence or presence of a morally important effect like a flood taking place—even though small amounts of GHGs add up to large amounts, and large amounts do make such differences.

To make plausible the metaphysical tolerance thesis about emissions and causation, Gesang invokes Nefsky's discussion of a voltmeter with a 1 kv margin of error. If after a series of single nanovolt increases the voltmeter registers '1 kv', Nefsky claims that we cannot say that the last nanovolt increase made the difference between its doing so and not doing so because its registering '1 kv' when it does is 'likely due to mechanical or environmental factors and not to the addition of some single nanovolt'.²⁹

However, without saying anything more about the mechanics of the device, this case appears to do nothing to support the claim that it is *impossible* for a single nanovolt to make a difference to whether the voltmeter registers '1 kv'. The most natural way for it to be likely that the last nanovolt failed to make the difference is for it to be likely that the voltmeter would have registered '1 kv' even if one had not added this nanovolt, since the other factors would have made it register '1 kv' anyway. But this certainly does not rule out there being some chance—perhaps quite small—that the single nanovolt increase did make the difference because the other factors would not have made the voltmeter register '1 kv' if one had not added the nanovolt.³⁰ The analogy to the climate case is thus simply what proponents of the expected consequences approach have been saying all along: for any climate harm, it is most likely that it will happen (or not) absent one's emission, but there is a small chance that one's emission will make the crucial difference and thus cause enormous harm.

²⁹ Nefsky (2012), p. 391.

³⁰ Nefsky actually acknowledges exactly this point on pp. 392–394 (*ibid.*).

But there could be a more interesting reason why we could not truly say that the last nanovolt made the difference to the voltmeter's registering '1 kv'. Suppose that the voltmeter operates in an objectively chancy or metaphysically indeterministic way, so that, given the entire history of the world at time t , the laws of nature fail to dictate that the voltmeter must be in some single state at $t+1$. Given this complete history, it is both possible for it to register '1 kv' and possible for it not to do so. As such, there will be no determinate fact of the matter about what would happen if one were to add or omit to add any given nanovolt. In some of the closest possible worlds in which one omits to add the nanovolt the voltmeter does not register '1 kv', but in others of these closest worlds the voltmeter does register '1 kv'. So for any single nanovolt increase followed by the voltmeter's registering '1 kv', it is not determinately true that this would not have happened had one not added the nanovolt—it might not have happened but it might still have happened.

Gesang's mention of quantum theory, the Copenhagen interpretation, and objective chances also support his being most charitably understood as proposing that the climate system may exhibit the foregoing sort of metaphysical indeterminacy.³¹ On this view, for small amounts of GHGs, there is simply no determinate fact of the matter about what would

31 Gesang (2017), p. 18. It is admittedly confusing that Gesang additionally mentions 'emergence theory', the content and relevance of which is unclear, and 'chaos theory', which Morgan-Knapp and Goodman convincingly argue is a plausible way to understand how small GHG emissions can determinately cause morally important effects—since it involves the climate system exhibiting sensitive dependence on initial conditions which gives rise to the above discussed butterfly effects by which events like small emissions can cascade into dramatic effects like floods and serious harms. It is also confusing that Gesang describes the voltmeter and climate scenarios as ones where the nanovolts and small GHG emissions determinately do not make any difference, as opposed to their simply being such that they do not determinately make a difference, because it is indeterminate whether they make a difference. But it seems to me that (i) there is no way to interpret Gesang as getting at a distinct convincing argument by invoking 'emergence' or 'chaos theory', and (ii) my interpretation of Gesang as getting at the above described kind of indeterminacy is the only way to interpret his remarks about the voltmeter and climate system in a way that (a) has him avoid simply being confused about the possibility of overdetermination not precluding the possibility of determinate difference making and the case being precisely the sort that Morgan-Knapp and Goodman were discussing, and (b) makes sense of his invocation of quantum theory and objective chances, which make perfect sense on an indeterministic interpretation of his remarks.

happen if we emit them. In some of the closest possible worlds in which we emit them, climate harms that otherwise might not occur do occur, while in others they do not occur. Our inability to be certain about what would happen stems not only from our ignorance of the exact details of the extremely complex climate system, but also from some of these details being undetermined by the laws of nature.

Gesang's objection to the expected consequences approach's account of the moral case for reducing our emissions is thus most charitably understood as turning upon the possibility of its being metaphysically indeterminate whether our acts will have certain outcomes. Nefsky offers a more general objection to the adequacy of the expected consequences approach that I think also turns upon a kind of indeterminacy, but indeterminacy in the value of the outcomes of our acts rather than whether our acts will cause them. Nefsky's criticism focuses on cases in which none of the possible outcomes of our acts seem to make a morally relevant difference. A classic example is Parfit's case of

The Harmless Torturers. Each of 1,000 torturers presses a button that increases by a tiny amount the electric current being fed to 1,000 victims. When none of the buttons are pressed the victims feel fine, but after each torturer has pressed her button each victim is being shocked by a massive current that causes her excruciating pain. Yet none of the victims seem to notice the additional electrical current from any one torturer's pressing her button.³²

Here, unlike in the cases discussed earlier, there do not seem to be any sharp thresholds that an additional button-pressing can cross to trigger a morally important outcome of pain or determinately worse pain to any victim. It thus seems that each torturer can be absolutely certain that because her pressing her button will not be noticed by any of the victims, it has no chance of making a morally relevant difference.

Kagan, Morgan-Knapp, and Goodman suggest that some cases of doing collective harm by polluting and causing climate change are *apparently imperceptible difference cases* of this kind.³³ But even if the most important environmental collective action problems are clear triggering

32 Parfit (1984), p. 80.

33 Kagan (2011), p. 129; Morgan-Knapp and Goodman (2015), p. 187.

cases, apparently imperceptible difference cases are still indirectly relevant to the expected consequences approach to explaining our moral reasons to reduce our ecological footprints. These cases challenge the general adequacy of the expected consequences approach, and it would seem problematic to accept this approach to our moral reasons in clear triggering cases but then to endorse a different explanation of our reasons in apparently imperceptible difference cases. Such a hybrid approach would among other things risk either (1) over-generating reasons in the clear triggering cases due to the reasons given by expected consequences combining with the other reasons operative in the apparently imperceptible difference cases, or (2) being implausibly ad hoc due to the lack of principled explanation as to why the non-expected-consequence-based reasons are present in the apparently imperceptible difference cases but not the clear triggering cases.

Kagan and several other proponents of the expected consequences approach have offered versions of the following 'reports-based argument' that apparently imperceptible difference cases are only apparent, and must actually involve sharp thresholds the crossing of which trigger determinately morally relevant differences. Suppose that the torturers press their buttons one after the other, and you ask a victim if she feels worse after each button-pressing. Because at the beginning she feels fine and at the end she feels awful, there must be some button-pressings after which she reports feeling worse. These button-pressings must have made a determinate difference to the badness of her pain.³⁴

But as Nefsky observes, if $S_0, \dots, S_{1,000}$ are the experiential states of a victim caused by 0 through 1,000 units of current affecting her, it seems entirely possible that she could report feeling worse in, say, S_{21} without S_{21} feeling determinately worse than S_{20} . Perhaps there was no clear point at which she noticed her states feeling determinately worse than S_0 , but this was somewhere around S_{15} - S_{25} , and it was somewhat random at which of these states she reported feeling worse to catch up to the fact that she seems to feel worse than she did in S_0 . Kagan actually concedes

34 Frank Arntzenius and David McCarthy, 'Self-torture and group beneficence', *Erkenntnis*, 47 (1997), 129–144 (pp. 132–135) <https://doi.org/10.1023/a:1005376607563>; Alastair Norcross, 'Comparing harms: Headaches and human lives', *Philosophy and Public Affairs*, 26 (1997), 135–167 (pp. 141–144), <https://doi.org/10.1111/j.1088-4963.1997.tb00079.x>; Kagan (2011), pp. 131–134; and Morgan-Knapp and Goodman (2015), pp. 186–190.

that our beliefs and reports about whether two very similar experiences were exactly the same are fallible indicators of whether they were.

Moreover, as Nefsky argues, proponents of the reports-based argument are too quick to dismiss the possibility of vague boundaries between which states in the series feel worse.³⁵ It might well be that for each pair of adjacent states S_i and S_{i+1} , there is no determinate fact of the matter about whether S_{i+1} feels worse to the victim than S_i . As Parfit has argued in another context, it does not seem that the degrees of badness of different painful experiences are precisely comparable. Consider a comparison between an intense pain that lasts for an hour and a much less intense pain that lasts longer. For sufficiently short durations like an hour and a minute, the less intense pain will be determinately less bad than the intense pain, and for sufficiently long durations like three years the less intense pain will be determinately worse. But it seems very implausible that there is some magic number of seconds or milliseconds of the less intense pain that makes it precisely as bad as the intense pain, such that one second or millisecond less would make it determinately less bad than the intense pain, and one second or millisecond more would make it determinately worse than the intense pain.³⁶

4. Indeterminacy and Supervaluationist Decision Principles

I have thus argued that Gesang's and Nefsky's challenges to the expected consequences approach hinge upon indeterminacy. Gesang has shown at most that there may be no determinate fact of the matter about whether our emissions will cause certain outcomes, while Nefsky has shown that there may be no determinate fact of the matter about whether certain of the outcomes caused by our acts are worse than their alternatives. But the expected consequences approach has been extended to these kinds of cases of indeterminacy. In this section I show how these extensions can be used to defend the approach against Gesang's and Nefsky's challenges.

³⁵ Nefsky (2012), pp. 380–387; Kagan (2011), p. 136.

³⁶ Derek Parfit, *On What Matters*, vol. 1 (New York: Oxford University Press, 2011), p. 132, <https://doi.org/10.1093/acprof:osobl/9780199572809.001.0001>.

If Gesang is correct that there is no fact of the matter about what would happen if we were to emit small amounts of GHGs (and we know this), then proponents of the expected consequences approach like Morgan-Knapp and Goodman are mistaken in claiming that there is a small *epistemic probability*—or degree of expectation we should have given our evidence—that such emissions will determinately trigger significant climate harms. But other proponents of the expected consequences approach have discussed how to apply it to cases in which it is indeterminate what the consequences of our acts will be. As Bykvist observes, if indeterministic processes yield precise objective chances of the outcomes that might obtain if we act in various ways,³⁷ we should follow

The Indeterministic Consequences Extension: apply the expected consequences approach using the objective chances as the probabilities of the outcomes to determine the (fact-relative) moral case for and against the different acts.³⁸

The rationale here seems exactly the same as that in favour of applying the expected consequences approach to cases where one's acts will have determinate but epistemically uncertain effects using the epistemic probabilities of those effects. When one cannot know with certainty

37 Where the objective chance of an outcome obtaining if an act is performed is something like the propensity of the act to lead to the outcome, or the proportion of the closest possible worlds in which the act is performed in which the outcome obtains. Such precise objective chances seem to be involved in the laws of nature according for instance to the Copenhagen interpretation of quantum mechanics—on which they determine such things as that a given electron has a 1/2 chance of being spin-up and a 1/2 chance of being spin-down.

38 Bykvist (2003), p. 30, n. 7; see also Douglas Portmore, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford: Oxford University Press, 2011), p. 56, n. 1, <https://doi.org/10.1017/s0953820812000490>. The fact-relative moral case for and against one's alternative acts is the extent to which the facts of one's circumstances count in favour of and against performing them, regardless of one's evidence about these facts (cf. Parfit (2011), pp. 150–162). The evidence-relative moral case is the moral case given one's evidence, which is typically at issue in discussions of the expected consequences approach. In a deterministic world the fact-relative moral case for and against an act is constituted by the benefits and harms it will actually bring about, regardless of one's evidence about this. But in an indeterministic world there are no determinate facts of the matter about what it will bring about: hence in such a world Bykvist suggests that we should use the expected consequences approach to determine what is the fact- and not simply the evidence-relative moral case for and against the act.

what would happen if one were to perform an act, in order to give appropriate weight to the ethical relevance of both

- i. the harms and benefits that might obtain if one were to perform it, and
- ii. the probabilities that these harms and benefits would obtain if one were to perform it,

one should weigh against and in favour of the act i in proportion to ii. It does not matter to the plausibility of this idea whether the lack of certainty and probabilities are due to one's ignorance of facts or the indeterministic structure of the world.

Even if we assume that our emissions do have precise objective chances of causing various climate harms, we presumably do not know exactly what these chances are. Given our evidence, all we seem to know is that an emission's objective chances of causing harms are proportional to its size, and we know of nothing else that increases or decreases these objective chances. As such, it seems that we should still follow Hiller and determine a given emission's epistemically expected harm by multiplying the total expected harms of our collective emissions by the ratio of the given emission to the collective amount. All Gesang's point about metaphysical indeterminacy does is entail that the epistemic probabilities of harm are generated not only by our ignorance of facts but also by the chancy nature of the world.

What, however, if Gesang's scenario of a metaphysically indeterministic climate system is one in which there are no precise objective chances of various outcomes occurring if various acts are performed? In such a scenario, while it may be determinate that the objective chance of a climate harm given some complete world history is greater than 0.5 and less than 0.9, there may simply be no fact of the matter about whether it is, say, 0.6531, 0.6527, or many other values between 0.5 and 0.9. While this cannot be handled by the indeterministic consequences extension alone, it can be handled by a natural and conservative extension of the expected consequences approach to cases of imprecise probabilities.

As Rinard has argued, there is a compelling way to motivate such an extension to cases in which the epistemic probabilities of various outcomes are imprecise. According to the general *supervaluationist*

approach to vagueness, for a claim with a vague predicate (such as ‘Singer is bald’) to be determinately true is for it to be true under every *admissible precisification*—or way of making the vague predicate precise (e.g. ‘someone with less than exactly 50% of typical hair-distribution is bald’) consistent with what is otherwise determinately true and false about the predicate (e.g. ‘someone with literally no hair is bald’ and ‘it is not the case that someone with a full typical hair-distribution is bald’).³⁹ If probabilities are imprecise, then claims about the probabilities of certain harms and benefits occurring if one performed certain acts seem to be clear instances of claims involving vague predicates. Thus, Rinard observes, the supervaluationist approach to vagueness supports

The Supervaluationist Principle for Probabilities: if probabilities are imprecise, but the expected harms of an act outweigh its expected benefits on every admissible precisification of the probabilities, then one determinately should not perform the act.

Suppose that, for every way of making precise the probabilities of the harms and benefits that might obtain if one were to perform an act, which is consistent with what one knows, the expected harms outweigh the expected benefits. Then the ethical significance of the act’s possible harms and their range of admissible likelihoods of occurring seems to unequivocally outweigh that of its possible benefits and their range of admissible likelihoods of occurring.⁴⁰ Note, moreover, that it does not matter to the plausibility of this rationale whether the imprecision of

39 Brian Weatherson, ‘The problem of the many’, in *The Stanford Encyclopedia of Philosophy*, ed. by Edward Zalta (2014), §7.3, <https://plato.stanford.edu/archives/win2016/entries/problem-of-many/>.

40 Rinard (2015), pp. 2–5. See also James Joyce, ‘A defense of imprecise credences in inference and decision making’, *Philosophical Perspectives*, 24 (2010), 281–323 (p. 311), <https://doi.org/10.1111/j.1520-8583.2010.00194.x>. As Joyce observes, what I am calling the supervaluationist principle for probabilities is a point of agreement among all proposed decision principles for the context of imprecise probabilities. This point of agreement is all I am relying upon in my argument. Rinard is distinctive in arguing that the general supervaluationist rationale supports not only the supervaluationist principle for probabilities but the further conclusion that, if expected consequences are different under different admissible precisifications of probabilities, then there is no determinate fact of the matter about what one should do. While I am sympathetic to this stronger view, my argument in no way depends upon it.

the probabilities is due simply to one's ignorance of facts that would justify having precise expectations, or to the world having an imprecise objectively chancy structure.

It seems, then, that if even if Gesang were right that a plausible climate scenario involves our emissions having objectively chancy effects without precise chances, we can use the supervaluationist principle for probabilities to apply the expected consequences approach. If we knew all there was to know about an act's objective chances of causing climate harms, we could apply the principle straightforwardly. For instance, if we knew that a given joyride had a propensity between one in one million and one in ten million of causing a flood that will inflict expected harm equal to the loss of 100 lives of 80 years, because the benefit to us is trivial in comparison to someone's losing anything between 70 and 7 days of life, the moral case against the joyride is determinately decisive.

Of course, if we assume any remotely realistic such scenario, we do not know all there is to know about our acts' imprecise objective chances of causing various climate harms. Given our evidence, all we seem to know is that an emission's admissible ranges of chances of causing harms are proportional to its size, and we know of nothing else that increases or decreases these ranges of objective chances. So it seems that, once again, we should use Hiller's method of determining an emission's epistemically expected harm by multiplying the total expected harms of our collective emissions by the ratio of the emission's size to that of the collective amount. The epistemic probabilities are generated not only by our ignorance of facts but by imprecise objective chances. But the epistemic probabilities of harm are exactly what they would be if they were, as Hiller, Morgan-Knapp, and Goodman suggest, due entirely to ignorance of facts.

Thus, even if Gesang is correct that our emissions may have imprecise objective chances of causing harm, modest extensions of the expected consequences approach entail that it is just as important to reduce our emissions as it would be if their effects were fully deterministic.

Rinard develops the supervaluationist approach to decision making in the context of indeterminate probabilities that certain outcomes will occur if we perform certain acts. But exactly parallel reasoning supports a similar approach to decision making in cases where the value of the outcomes of our acts is indeterminate. As Nefsky suggests, apparently

imperceptible difference cases such as that of Parfit's Harmless Torturers seem to be cases of this kind. Pressing a button that marginally increases the amount of current flowing to the 1,000 victims will, for at least some ways the other torturers might act, neither determinately worsen the pain of any victims nor determinately leave them no worse off than they otherwise would have been.

Just as the supervaluationist approach to vagueness supports the supervaluationist principle for probabilities, it also supports

The Supervaluationist Principle for Values: if the degrees of harm or benefit of some possible outcomes of our acts are imprecise, but the expected harms of an act outweigh its expected benefits on every admissible precisification of these degrees of harm or benefit, then one determinately should not perform the act.

Much as above, suppose that, for every way of making precise the degrees of harm or benefit of the possible outcomes of an act that are consistent with what we know, its expected harms outweigh its expected benefits. Then the ethical significance of the range of the act's admissible degrees of harm and their probabilities of occurring seems to unequivocally outweigh that of the range of the act's admissible degrees of benefit and their probabilities of occurring.

To apply this to the Harmless Torturers case, again let $S_0, \dots, S_{1,000}$ be the experiential states of a victim caused by 0 through 1,000 units of current. For at least some x , it is not determinate whether S_{x+1} is worse than S_x , but $S_{1,000}$ is determinately much worse than S_0 . What we must do is consider the set of all admissible precisifications or ways of assigning harm or disvalue to S_1 through $S_{1,000}$ consistent with $S_{1,000}$ being much worse than S_0 . Each such precisification distributes the difference in harm between $S_{1,000}$ and S_0 , $H(S_{1,000}) - H(S_0)$, among the various states. So on each admissible precisification there will be n ($\leq 1,000$) states with amounts of additional harm h_1, \dots, h_n of $H(S_{1,000}) - H(S_0)$, such that $h_1 + \dots + h_n = H(S_{1,000}) - H(S_0)$. By pressing one's button one has an equal chance of causing any of S_1 through $S_{1,000}$, and the amount of additional harm of any state not assigned an amount of additional harm by an admissible precisification is zero. So on each admissible precisification, one's expected harm from pressing one's button will

be $(1/1,000) \times h_1 + \dots + (1/1,000) \times h_n = (h_1 + \dots + h_n) / 1000 = (H(S_{1,000}) - H(S_0)) / 1,000$, which is equivalent to that of a 1/1,000 chance of causing the full difference in harm between $S_{1,000}$ and S_0 . Since one is doing this to 1,000 victims, the expected harm one is causing to all of them is equal to that of taking a single victim from S_0 to $S_{1,000}$ with certainty.

But this is the exact same expected harm from pushing the button that we would get if we followed the advocates of the reports-based argument who insist that it cannot be vague whether certain states feel worse than others in apparently imperceptible difference cases. These authors simply insist that, for each victim there is some single actual distribution of the difference in harm between $S_{1,000}$ and S_0 among the states S_1 through $S_{1,000}$, $\hat{h}_1, \dots, \hat{h}_n$, such that $\hat{h}_1 + \dots + \hat{h}_n = H(S_{1,000}) - H(S_0)$ (which they think corresponds in some way to what she would report if she were asked).

Thus, even if (as I suspect) Nefsky is correct that there are vague boundaries among the harms caused to the victims in states S_1 through $S_{1,000}$, on every admissible precisification of the harms, the expected harms to the victims will be identical to what they would be if there were no such vague boundaries. So, given the supervaluationist principle for values, this means that the moral case against pressing the button will be decisive when the differences in harm are vague just in case it would also be decisive if these differences were not vague. So once again, even if there is indeterminacy where advocates of the expected consequences approach have thought that there is determinacy, our modest extension of the expected consequences approach enables it to yield the exact same results that it would if there was actually determinacy.

5. Conclusion

The expected consequences approach provides a clear explanation of why we should take both political and personal action to reduce our ecological footprints. In environmental collective action problems like that of anthropogenic climate change, the collective result of our acts like emitting GHGs is extremely harmful. As Hiller, Broome, Morgan-Knapp, and Goodman have argued, emissions far less than the sum-total of all anthropogenic GHGs have small chances of crossing thresholds and causing the more likely butterfly effects of the chaotic climate system that result in great harm. Because an emission's chance of having

these effects is proportional to its size, and we know of nothing else that makes any given emission more or less likely to have them, we can determine an emission's expected harms by multiplying the total expected harm of our collective emissions by the ratio of the emission's size to that of the collective amount.

Following Nolt, we can use this method to estimate the expected benefit of various political and private acts of reducing our destructive impacts on the environment. Doing so, we can see that policies like a Green New Deal in the US could reduce the emissions of CO₂ equivalent by hundreds of billions of tonnes, which can be expected to save hundreds of millions of lives. As such, we must vote and campaign for policies like the Green New Deal because, although our votes and those of others we may convince are very likely to make no difference to the enactment of such policies, they have a small chance of winning the election and making all the difference. This makes the expected benefit of an average vote in such an election equivalent to that of saving between 5 and 10 full human lives, and that of even a vote in an extremely uncompetitive state equivalent to extending someone's life by almost a year, which decisively outweighs the costs to us of voting and canvassing.

This applies just as much to actions within our personal as our political lives. A single choice of purchasing plant-based alternatives rather than a chicken's corpse has a small chance of omitting to cause an enormous number of additional chickens to be tortured and killed in a way that would involve much greater ecological destruction than the production of plant-based alternatives. This makes the expected benefit of a single choice of purchasing plant-based alternatives rather than a chicken corpse equivalent to that of omitting to torture and kill one chicken and omitting the ecological damage of torturing and killing her, which decisively outweighs the relatively trivial cost of forgoing familiar taste e-pleasures from eating her corpse. By going vegan one can be expected each year to directly prevent the torture and killing of somewhere between at least 232 and 443 vertebrate animals, together with the benefits to wild animals and other humans of using up to 2.83 fewer acres or only 1/18 as much land, and emitting at least 1.5 fewer tonnes of CO₂ equivalent.⁴¹ We can also prevent a great deal of

41 Oppenlander (2013); Harish Sethu, 'How many animals does a vegetarian save?', *Counting Animals* (February 6, 2012), <http://www.countinganimals.com/how-many-animals-does-a-vegetarian-save/>.

expected harm by reducing our flying and using video technology to hold our academic meetings online without the need to travel. A single econference like *Around the World* can be expected to save 300 tonnes of CO₂ equivalent, which prevents an expected harm comparable to that of killing a 69-year-old Canadian, his 14 elderly dogs and the 1,400 elderly fish in a nearby pond.

The adequacy of this account has been challenged by Gesang's contention that it may be not only unlikely but genuinely impossible for sufficiently small emissions to make a morally relevant difference, and Nefsky's concern that the expected consequences approach cannot explain why we should avoid contributing to collective harm in cases like the *Harmless Torturers* where the possible outcomes of our acts do not seem determinately worse for anyone.

In response I have argued that Gesang has shown at most that, if the laws of nature are not deterministic, there may be no determinate fact of the matter about what would happen if we emit some quantity of GHGs. But, as I have explained, the expected consequences approach has already been extended to such cases of indeterminism. If indeterministic outcomes occur with precise objective chances, the general idea of the expected consequences approach supports using those objective chances—or our best estimates of them—in our calculations of expected consequences. If the objective chances are not precise, then we should follow the supervaluationist approach that has already been developed for imprecise epistemic probabilities, and conclude that, if on every admissible precisification of the probabilities the expected harm of emitting outweighs its expected benefits, then we should not emit. From our evidential perspective we do not actually know the precise or imprecise objective chances of any given emission making any given difference—all we seem to know is that the chance or range of chances of the emission causing harms is proportional to its size. So we should still determine an emission's epistemically expected harm by multiplying the expected harm of our collective emissions by the ratio of the emission's size to that of the collective amount. This means that even if Gesang is correct that our emissions may have imprecise objective chances of causing harm, the expected consequences approach entails that it is just as important to reduce them as it would be if their effects were fully deterministic.

Exactly similar reasoning can be used to respond to Nefsky's concern that the outcomes of our acts in cases like the Harmless Torturers may not be determinately worse for the victims. The supervaluationist approach here supports the conclusion that, if it is vague which mental states of the victims are worse for them than others, then we should not perform the act if its expected harms outweigh its expected benefits on every admissible precisification of the harms to the victims. But every admissible precisification must distribute the full difference between no one contributing to the victims' harm and everyone contributing to it among the various amounts of positive contribution, each of which our contribution has an equal chance of causing. So on every admissible precisification, the expected harm is equal to the total difference in harm divided by the number of states of positive contribution. But this is exactly what the expected harm would be if there were one actual precise distribution of the harm among the states of positive contribution. So, if there is indeterminacy in the value of the outcomes of our acts—just as if there is indeterminacy in the probabilities of certain outcomes of our acts occurring—the expected consequences approach entails that it is just as important to omit contributing to the victims' harm as it would be if there was no such indeterminacy.

I therefore conclude that the expected consequences approach can surmount the challenges of indeterminacy posed by Gesang and Nefsky, and that it provides a fully adequate account of our powerful moral reasons to take both political and personal action to reduce our ecological footprints.

Bibliography

- Academy of Nutrition and Dietetics, 'Position of the Academy of Nutrition and Dietetics: Vegetarian diets', *Journal of the Academy of Nutrition and Dietetics*, 116 (2016), 1970–1980, <https://doi.org/10.1016/j.jand.2016.09.025>
- Arntzenius, Frank, and David McCarthy, 'Self-torture and group beneficence', *Erkenntnis*, 47 (1997), 129–144, <https://doi.org/10.1023/a:1005376607563>
- Balcombe, Jonathan, *Pleasurable Kingdom: Animals and the Nature of Feeling Good* (London: Macmillan, 2006).

- Böhm-Bawerk, Eugen, *The Positive Theory of Capital*, trans. by William Smart (1889, reprinted in New York: G. E. Stechert & Co., 1930), <https://doi.org/10.2307/2956155>
- Braithwaite, Victoria, *Do Fish Feel Pain?* (New York: Oxford University Press, 2010), <https://doi.org/10.1086/656881>
- Broome, John, *Climate Matters: Ethics in a Warming World* (New York: W. W. Norton & Co., Inc., 2012).
- Bykvist, Krister, 'Normative supervenience and consequentialism', *Utilitas*, 15 (2003), 27–49, <https://doi.org/10.1017/s0953820800003757>
- Chow-Fraser, Trevor, Chelsea Miya and Oliver Rossier, *Moving Ideas without Moving People: How to Econference at the University of Alberta* (2018), <https://aroundtheworld.ualberta.ca/wp-content/uploads/2018/05/E-Conferencing-Toolkit.pdf>
- Edelman, David, and Anil Seth, 'Animal consciousness: A synthetic approach', *Trends in Neurosciences*, 32 (2009), 476–484, <https://doi.org/10.1016/j.tins.2009.05.008>
- Eshel, Gidon, and Pamela Martin, 'Diet, energy, and global warming', *Earth Interactions*, 10 (2006), 1–17, <https://doi.org/10.1175/ei167.1>
- Food and Agriculture Organization of the United Nations, *Tackling Climate Change Through Livestock: A Global Assessment of Emissions and Mitigation Opportunities* (Rome: FAO, 2013), <http://www.fao.org/3/i3437e/i3437e00.htm>
- Flynn, Mary, and Andrew Schiff, 'Economical healthy diets (2012) including lean animal protein costs more than using extra virgin olive oil', *Journal of Hunger & Environmental Nutrition*, 10 (2015), 467–482, <https://doi.org/10.1080/19320248.2015.1045675>
- Gelman, Andrew, Nate Silver and Aaron Edlin, 'What is the probability that your vote will make a difference?', *Economic Inquiry*, 50 (2012), 321–326, <https://doi.org/10.1111/j.1465-7295.2010.00272.x>
- Gesang, Bernward, 'Climate change—do I make a difference?', *Environmental Ethics*, 39 (2017), 3–19, <https://doi.org/10.5840/enviroethics20179261>
- Goodland, Robert, and Jeff Anhang, 'Livestock and climate change: What if the key actors in climate change are...cows, pigs and chickens?', *WorldWatch*, November/December 2009, <https://awellfedworld.org/wp-content/uploads/Livestock-Climate-Change-Anhang-Goodland.pdf>
- Hiller, Avram, 'Climate change and individual responsibility', *The Monist*, 94 (2011), 349–368, <https://doi.org/10.5840/monist201194318>
- Hiller, Avram, 'Morally significant effects of ordinary individual actions', *Ethics, Policy and Environment*, 14 (2011), 19–21, <https://doi.org/10.1080/21550085.2011.561588>

- Hiltner, Ken, 'A nearly carbon-neutral conference model', <https://hiltner.english.ucsb.edu/index.php/ncnc-guide/>
- IPCC 2014, *Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, ed. by Core Writing Team, R. K. Pachauri and L.A. Meyer (Geneva: Switzerland, 2014), https://www.ipcc.ch/site/assets/uploads/2018/05/SYR_AR5_FINAL_full_wcover.pdf
- Jackson, Frank, and Michael Smith, 'Absolutist moral theories and uncertainty', *The Journal of Philosophy*, 103 (2006), 267–283, <https://doi.org/10.5840/jphil2006103614>
- Joyce, James, 'A defense of imprecise credences in inference and decision making', *Philosophical Perspectives*, 24 (2010), 281–323, <https://doi.org/10.1111/j.1520-8583.2010.00194.x>
- Kagan, Shelly, 'Do I make a difference?', *Philosophy and Public Affairs*, 39 (2011), 105–141, <https://doi.org/10.1093/0198239165.001.0001>
- Kagan, Shelly, *The Limits of Morality* (Oxford: Oxford University Press, 1989), <https://doi.org/10.1111/j.1088-4963.2011.01203.x>
- Kremp, Pierre-Antoine, and Andrew Gelman, 'What is the chance that your vote will decide the election?', https://pkremp.github.io/pr_decisive_vote.html
- 'List of close election results', *Wikipedia, The Free Encyclopedia* (October 9, 2020), https://en.wikipedia.org/wiki/List_of_close_election_results
- Mann, Michael, *The Hockey Stick and the Climate Wars: Dispatches from the Front Lines* (New York: Columbia University Press, 2012), <https://doi.org/10.7312/mann15254>
- Matheny, Gaverick, 'Expected utility, contributory causation, and vegetarianism', *Journal of Applied Philosophy*, 19 (2002), 293–297, <https://doi.org/10.1111/1468-5930.00223>
- Morgan-Knapp, Christopher, and Charles Goodman, 'Consequentialism, climate harm, and individual obligations', *Ethical Theory and Moral Practice*, 18 (2015), 177–190, <https://doi.org/10.1007/s10677-014-9517-9>
- National Highway Traffic Safety Administration, *Draft Environmental Impact Statement for The Safer Affordable Fuel-Efficient (SAFE) Vehicles Rule for Model Year 2021–2026 Passenger Cars and Light Trucks* (2018), https://www.nhtsa.gov/sites/nhtsa.dot.gov/files/documents/ld_cafe_my2021-26_deis_0.pdf
- Nefsky, Julia, 'Consequentialism and the problem of collective harm: A reply to Kagan', *Philosophy and Public Affairs*, 39 (2012), 364–395, <https://doi.org/10.1111/j.1088-4963.2012.01209.x>
- Nolt, John, 'How harmful are the average American's greenhouse gas emissions?', *Ethics, Policy and Environment*, 14 (2011), 3–10, <https://doi.org/10.1080/21550085.2011.561584>

- Norcross, Alastair, 'Comparing harms: Headaches and human lives', *Philosophy and Public Affairs*, 26 (1997), 135–167, <https://doi.org/10.1111/j.1088-4963.1997.tb00079.x>
- Norcross, Alastair, 'Puppies, pigs, and people: Eating meat and marginal cases', *Philosophical Perspectives*, 18 (2004), 229–245, <https://doi.org/10.1111/j.1520-8583.2004.00027.x>
- Nye, Howard. 'Chaos and Constraints,' in *Dimensions of Moral Agency*, ed. By David Boersema (Newcastle: Cambridge Scholars Publishing, 2014), pp. 14–29, <https://www.cambridgescholars.com/resources/pdfs/978-1-4438-6692-7-sample.pdf>
- Oppenlander, Richard, *Food Choice and Sustainability: Why Buying Local, Eating Less Meat, and Taking Baby Steps Won't Work* (Minneapolis: Langdon Street Press, 2013).
- Panksepp, Jaak, *Affective Neuroscience: The Foundations of Human and Animal Emotions* (Oxford: Oxford University Press, 1998), <https://doi.org/10.1017/s0317167100052070>
- Parfit, Derek, *Reasons and Persons* (New York: Oxford University Press, 1984), <https://doi.org/10.1093/019824908x.001.0001>
- Parfit, Derek, *On What Matters*, vol. 1 (New York: Oxford University Press, 2011), <https://doi.org/10.1093/acprof:osobl/9780199572809.001.0001>
- Poore, Joseph and Thomas Nemcek, 'Reducing food's environmental impacts through producers and Consumers', *Science*, 360 (2018), 987–992, <https://doi.org/10.1126/science.aaq0216>
- Portmore, Douglas, *Commonsense Consequentialism: Wherein Morality Meets Rationality* (Oxford: Oxford University Press, 2011), <https://doi.org/10.1093/acprof:oso/9780199794539.003.0007>
- Ramsey, Frank P., 'Truth and probability' (1926), reprinted in *Decision, Probability, and Utility: Selected Readings*, ed. by Peter Gärdenfors and Nils-Eric Sahlin (New York: Cambridge University Press, 1988), pp. 19–47, <https://doi.org/10.1017/cbo9780511609220>
- Rinard, Susanna, 'A decision theory for imprecise probabilities', *Philosophers' Imprint*, 15 (2015), 1–16, <http://hdl.handle.net/2027/spo.3521354.0015.007>
- Roberts, Melinda, 'The asymmetry: A solution', *Theoria*, 77 (2011), 333–367, <https://doi.org/10.1111/j.1755-2567.2011.01117.x>
- Ross, W. D., *The Right and the Good* (Oxford: Clarendon Press, 1930).
- Rothbard, Murray, *Man, Economy, and the State* (Princeton: Van Nostrand, 1962; reprinted in Auburn: Ludwig von Mises Institute, 2001).
- Savage, Leonard J., *The Foundations of Statistics* (New York: Dover Publications, Inc., 1954).

- Searchinger, Timothy, Stefan Wirseniuss, Tim Beringer, and Patrice Dumas, 'Assessing the Efficiency of Changes in Land Use for Mitigating Climate Change', *Nature* 564 (2018), 249–253, <https://doi.org/10.1038/s41586-018-0757-z>
- Sethu, Harish, 'How many animals does a vegetarian save?', *Counting Animals* (February 6, 2012), <http://www.countinganimals.com/how-many-animals-does-a-vegetarian-save/>
- Shepon, Alon, Gideon Eshel, Elon Noor, and Ron Milo, 'The Opportunity Cost of Animal Based Diets Exceeds All Food Losses', *Proceedings of the National Academy of Science* [PNAS] 115(15) (2018), 3804–3809, <https://doi.org/10.1073/pnas.1713820115>
- Shiffrin, Seana, 'Wrongful life, procreative responsibility, and the significance of harm', *Legal Theory*, 5 (1999), 117–148, <https://doi.org/10.1017/s1352325299052015>
- Sidgwick, Henry, *The Methods of Ethics*, 7th edn (London: Macmillan and Co., Limited, 1907).
- Singer, Peter, 'Utilitarianism and vegetarianism', *Philosophy and Public Affairs*, 9 (1980), 325–337.
- Singer, Peter, and Jim Mason, *The Ethics of What We Eat: Why Our Food Choices Matter* (New York: Rodale Press, 2006).
- Sinnott-Armstrong, Walter, 'It's not my fault: Global warming and individual moral obligations', in *Perspectives on Climate Change: Science, Economics, Politics, Ethics*, ed. by Walter Sinnott-Armstrong and Richard Howarth (Amsterdam: Elsevier, 2005), pp. 285–307, [https://doi.org/10.1016/S1569-3740\(05\)05013-3](https://doi.org/10.1016/S1569-3740(05)05013-3)
- Tillett, Emily, 'Virginia election results 2017: Republican David Yancey wins Virginia House seat', *CBS News* (January 4, 2018), <https://www.cbsnews.com/news/virginia-election-results-lottery-drawing-house-of-delegates-david-yancy-winner-virginia-house-seat/>
- Tomasik, Brian, 'How many wild animals are there?', 2018, <http://reducing-suffering.org/how-many-wild-animals-are-there>
- von Neumann, John, and Oskar Morgenstern, *Theory of Games and Economic Behavior* (Princeton: Princeton University Press, 1944), <https://doi.org/10.1515/9781400829460>
- Weatherston, Brian, 'The problem of the many', in *The Stanford Encyclopedia of Philosophy*, ed. by Edward Zalta (2014), <https://plato.stanford.edu/archives/win2016/entries/problem-of-many/>
- Wolf, Julie, Ghassem Asrar and Tristram West, 'Revised methane emissions factors and spatially distributed annual carbon fluxes for global livestock', *Carbon Balance and Management*, 12 (2017), 1–24, <https://doi.org/10.1186/s13021-017-0084-y>