Utilitarianism and Heuristics

Bernward Gesang[1]

1 Methodic and Empirical Arguments against Expected Utility Analyses

W. Sinnott-Armstrong states that individuals have no moral obligation to avoid unnecessary drives with regard to global warming, which "will still occur even if I do not drive just for fun."[1] This constitutes a provocation. A consequentialist ethicist is obligated to examine this statement by calculating the expected utility with regard to a single car ride. The expected utility is a product of the scale of the utility (including harm) along with the likelihood that these effects will occur. Such calculations of expected utility are the standard consequentialist response to the problem as to how the consequences of an action can be calculated in advance: "But if my discussion of this sort of case (call it a 'relevance of a single action type of case,' B.G.) is correct, then the consequentialist can handle such cases using the familiar appeal to expected utility."[2] The consequentialist argues that he is capable of orienting his decisions towards this sort of knowledge about consequences. The criticism has always been made that such calculations of expected utility are too demanding in an epistemic sense and are overly elaborate in their actual execution.

Let us take the statement by Sinnott-Armstrong. What do we need to know in order to decide whether he is correct? A. Hiller calls upon us to make the following calculation:

"(Step 1) Estimate the amount of GHG emitted by the one drive d.

---

[1] Walter Sinnott-Armstrong, "It's Not My Fault: Global Warming and Individual Moral Obligations," *Perspectives on Climate Change: Science, Economics, Politics and Ethics, Advances in the Economics of Environmental Research* 5 (2005): 285–307, p. 297.

[2] Shelly Kagan, "Do I make a difference?" *Philosophy & Public Affairs* 39 (2011): 105–141, p. 129.

✉ Bernward Gesang
    Bernward.Gesang@uni-mannheim.de

[1]   Philosophisches Seminar, Universität Mannheim, Mannheim, Germany

Ⓐ Springer

(Step 2) Estimate the total amount of GHG emissions responsible for climate change e.

(Step 3) Estimate the total amount of harm that climate change will cause h.

(Step 4) Calculate (d/e) x h."[3]

This is not a matter of an expected-utility analysis because no attention is paid to the issue of probable occurrence of effects. In this calculation of costs and utility, the sole focus is on the factor of the extent of harm. But here as well, there are multiple grounds for criticism: S. Kagan criticized this kind of calculation as "bookkeeping" because of its statistical allocation of harm to the individual.[4] The consequentialist, who in this case is a utilitarian, wants to know whether his concrete car ride makes a *causal difference* for the world. But I have further reservations with regard to the pure feasibility of the calculatory steps: Already Step 1 disregards the emissions arising during the production of automobiles; i.e. a complete supply chain analysis is lacking. I consider the estimates called for in Step 3 to be nothing other than illusory: For example, C. Lumer tried to do this, and the attempt failed with regard to the empirical prognoses, because the rise of China was not anticipated.[5] It is possible to discuss whether this constitutes a failure of the method. I believe that these problems demonstrate that the data for so complex an analysis cannot be sufficiently determined. One neither knows nor can estimate how much harm will arise through climate change. That is within the purview of omniscient beings alone, because the amount of data and of possible alternatives and interconnections is unmanageably complex. So the introductory statement by W. Sinnott-Armstrong simply cannot be evaluated, even if we exclude the factor of the probable occurrence of effects.

Expected utility analyses as means of decision have a purely procedural complexity, which makes them subject to error; F. Feldman points to this shortcoming with what tends to be a methodic argument:

"But the determination of expected utility is even more problematic. To perform the relevant calculations, one needs to know his alternatives, and for each alternative, one also needs to know every one of its possible outcomes, the actual value of each of those outcomes and the probability of each outcome, given the act, on one's current evidence. One also needs to know some mathematical facts: the product, for each outcome, of its value and its probability, and the sum of those products. If the epistemic task in the actual utility case was daunting, this task is double-daunting." [6] Feldman concludes from this that, if impracticality is the problem of consequentialist analyses, a calculation of expected utility is not the solution.[7]

---

[3] Avram Hiller, "Climate Change and Individual Responsibility," *The Monist* 94 (2011): 349–368, p. 357.

[4] See Jonathan Glover, Martin J. Scott-Taggart, "It Makes No Difference Whether or Not I Do It," *Aristotelian Society Supplementary* 49 (1975): 171–210, p. 174.

[5] See Christoph Lumer, *The Greenhouse* (Lanham et al.: University Press of America, 2002), pp. 71–75.

[6] Fred Feldman, "Actual Utility, the Objection from Impracticality, and the Move to Expected Utility," *Philosophical Studies* 129 (2006): 49–79, p. 56.

[7] In expected utility calculations, *metrical distances* are calculated between outcomes—with the use, however, of subjective probabilities. See Feldman, op.cit., p. 62.

Because of these problems, decision theory includes entire theoretical branches which choose to do without any sort of probabilities. As G. Betz describes: "Rather than building on probabilistic guesswork, we should acknowledge the full extent of our ignorance and the uncertainty we face. We should not simply make up the numbers. And we should refrain from wishful thinking."[8]

R. Hare's "two-level-morality," inspired by J.S. Mill, provides us with an instrument for solving problems of decisions about moral problems.[9] Hare offers a heuristic for decision-making with regard to unproblematic everyday situations in which a selection was made of intuitive, utility-generating principles that do not contradict each other.[10] An *intuitive level* where simple criteria based on everyday morality are used for making decisions is complemented by a *critical level* where calculations of utility are undertaken when, for example, intuitive criteria contradict each other. This heuristics was established in order to get a grip on the aforementioned problems. In a normal case, conflicts are resolved on the intuitive level; it is only in a special case that the critical level must be included. This prevents continuous calculation and the ensuing complexities. I believe, however, that on the intuitive level the principles "act efficiently" and "try to make a difference" would also have to be selected and that, among other things, they give rise to contradictions with other principles. There are accordingly a certain number of cases which must be resolved on the second, critical or act-utilitarian level in which relevant calculations must be made.[11] But even in a two-level morality such as that of Hare or Mill, it is difficult to achieve a determination of expected utility. Hare remains silent regarding how calculations should be made on the critical level; he can only claim that the number of affected cases is decreasing. As Feldman says, it requires a significant knowledge of possible actions and their results, as well as of estimates of utility values and probabilities. While not impossible to carry out, this is complicated and, unfortunately, also sometimes imprecise and uncertain. In fields of application such as climate ethics, in which there is a high degree of uncertainty, these problems often make it impossible to arrive at a result[12]. Hence I believe that we need heuristics which, extending past level one, identify particularly problematic calculations on level two *and also offer assistance in this regard*.

This essay will first demonstrate that there is a methodic and an empirical challenge for the thesis of consequentialism that actions can be measured by their results. In a second step, it will become clear that this challenge can be extended to apply to

---

[8] Gregor Betz, "Accounting for Possibilities in Decision Making," in S. Hansson, G. Hirsch Hadorn (eds.), *The Argumentative Turn in Policy Analysis* (Cham (Switzerland): Springer, 2016): 135–169, p. 137.

[9] See Richard M. Hare, *Moral Thinking: Its Levels, Method and Point* (Oxford: Oxford University Press, 1981) and John S. Mill, *Utilitarianism*, ed. A. Ryan (London et al.: Penguin 1987), 272–338, pp. 40–44.

[10] The automatic choices which are based on virtuous attitudes also belong to this level. See Christopher Woodard, *Taking Utilitarianism Seriously* (Oxford: Oxford University Press, 2019), chapter 9.

[11] With Hare, critical thought selects the principles on the intuitive level (Hare 1981, ch. 1.3.3) in contrast to Mill, who accepts intuitive morality on this level: i.e. Hare has more intensive need of the second level.

[12] Bernward Gesang "Climate Change: Do I make a difference?" in: *Environmental Ethics* Vol. 39, 2017, 3–19.

collective actions problems as well. In fields of great uncertainty, an expected utility analysis as a decision tool[13] frequently imposes an excessive epistemic demand on consequentialism. This essay will accordingly seek heuristics which support the decision-making of utilitarians and which reduce uncertainty especially with regard to what constitutes correct behavior in collective action problems. The answer formulated in response to these problems will involve a concentration on the realization of comparatively large potential utility (i.e. great events) instead of on symbolic collaborative projects. Another example for the problem of outcome prognoses is offered by the debate over a *cost-benefit analyses.* These are made in order to calculate, for example, the effects of large construction projects such as the Jangtsekiang Dam in China on social welfare. In this case, calculations are made for values which are considered to be incomparable, such as economic growth and impact on archaeological sites, through methods such as "hedonic pricing." Here, the debates as well as the complexity of the undertaking (calculation of all relevant, often hidden consequences, comparative methods for the determination of a common denominator, consideration of the subjective significance of money) show that the reproof of subjectivity may justifiably be made with regard to the project.[14] This is evident, for example, in the regular occurrence of different results when analyses are made of the same subject.[15] The question presents itself as to what consequences this has for the methodology of ethical decision-making: a) When do determinations of expected utility become particularly prone to error? b) How can we transform the criterion of consequentialism regarding ethical correctness into a decision criterion that is not based on doubtful calculations?

The fact that prognoses are often undependable is likewise shown by an empirical reservation regarding the dependability of prognoses which are based on the *estimates of experts*. D. Kahneman argues firstly against the reliability of prognoses in general and observes that our spontaneous thought (system one) works inductively, develops active, coherent narratives and thereby distorts facts. It gives us the illusion of understanding the past, whereupon we erroneously believe that we can also predict the future.[16] We tend to overestimate our prognostic capabilities.

Secondly, this diagnosis applies not only to spontaneous thought. As a rule, expert opinions are also faulty, at least with regard to long-term prognoses. Kahneman speaks of a competency illusion: "People who spend their time, and earn their living, studying a particular topic produce poorer predictions than dart-throwing

---

[13] I do not consider expected utility as a full replacement for data on objective utility. However, it can be a means that helps us to estimate objective utility. See Graham Oddie, Peter Menzies, „An Objectivist´s Guide to Subjective Value," in: P. Pettit (ed.), *Consequentialism* (Aldershot/Brookfield: Dartmouth, 1993), 327–350. When I refer to this means, I explicitly speak of "expected utility".

[14] See Jean Dréze, Nicholas Stern (1987), "The cost-benefit analysis," in M. Feldstein, A.J. Auerbach (eds.), *Handbook of Public Economics 2* (North Holland Imprint): 909–989, pp. 955–960.

[15] Mark Sagoff, *Price, Principle, and the Environment* (Cambridge: Cambridge University Press, 2004), p. 109; Fearnside, Philip M. (2012) *The canadian feasibility study of the three gorges dam proposed for china's yangzi river : a grave embarrassment to the impact assessment profession*, 9165. https://doi.org/10.1080/07349165.1994.9725849.

[16] See Daniel Kahneman, *Thinking, Fast and Slow* (London et al.: Farrar, Straus and Giroux, 2011), p. 201.

monkeys who would have distributed their choices evenly over the options."[17] Short-term prognoses are more likely to be predictable, but do "not expect much from pundits making long-term forecasts."[18] Formulas and algorithms are more reliable, but they as well quickly reach their limits in the case of long-term predictions.[19] Must we refrain from making long-term predictions altogether, regardless of whether they are of a spontaneous or scientific nature, because the world is inherently unpredictable?[20] If that were true, then we could refrain from a large part of politics and from our actions on a social level. This degree of strict consistency seems to be excessive. But plausible reservations regarding determinations of expected utility have also emerged in the discussion regarding Kahneman. Attention should be paid to these reservations.

But where do our decision-making calculations fail in particular? What are we most likely to know regarding the morally relevant differences among the consequences of actions; how can relevant information be extended into heuristics which are helpful especially with collective action problems?[21]

2 Collective Action as Uncertain Action

In recent years, there has been a detailed discussion of collective action problems in the literature.[22] What is the focus here? If a large enough number of people act in a certain manner (something which would have to be described in more detail), the results are quite negative*; but none* of the individual actions constituting the mode of bad behavior seem to make a difference with regard to the outcome. Thus

---

[17] Ibid., p. 219.

[18] Ibid., p. 221.

[19] See Ibid., p. 227.

[20] It can of course be the case that in the long term, we cannot predict anything; that would conform to a thesis often proposed by Kahneman. This quite insistent thesis, however, stands on wobbly legs. Kahneman's main source for this sort of skepticism regarding prognoses, namely a book by Tetlock, is far less skeptical. Tetlock proposes consulting self-critical experts who do not suffer from hubris. See Philip E. Tetlock, *Expert Political Judgement: How Good is it? How Can We Know?* (Princeton: Princeton University Press, 2005), p. 21.

He distances himself from a radical skepticism regarding prognoses and calls for a "skeptical meliorism" which would include the possibility of self-critical predictions. See Tetlock, op. cit., p. 19. Furthermore, *Kahneman himself does not maintain the skeptical attitude he propagates.* After he has professed skepticism, he speculates in the next chapter of his book about long-term prognoses concerning how costs develop in large construction projects and proposes freeing them from "fallacies in planning." See Kahneman, op. cit., p. 250. But significant long-term prognoses apparently are possible if they are carefully prepared, even if the epistemic uncertainty undoubtedly gives us pause for thought. See Kahneman, op. cit., pp. 220, 241.

[21] This list of heuristics discussed in this essay is not complete. For example, Betz presents further procedures. See Betz, op. cit.

[22] I will focus here on "big" cooperative actions whose consequences are difficult to predict. When only five protagonists are involved in a collaboration, I can more easily predict its outcome than when millions of persons are participating. See Glover/Scott-Taggart, op. cit., Kagan, op. cit., Julia Nefsky, "Consequentialism and the problem of collective harm: A reply to Kagan," *Philosophy & Public Affairs* 39 (2012): 364–395,

candidates for these sorts of problematic manners of cooperative action consist of actions whose individual contribution to an overall outcome is extremely small or scarcely perceptible. An example is climate protection, where the total amount of trace-gas emissions has a disastrous effect, but there is controversy as to whether a single car trip by one protagonist makes a difference that is morally relevant or causally identifiable.[23] Other problems of collective action may be found, for example, in consumer ethics when it is a matter of whether a single purchase of a "good product" from mass production (e.g. a T-shirt) can make a difference. The problem with such cases is that suffering is created by an accumulation of human actions while it remains unclear as to whether any particular protagonist per se should have acted differently.

I doubt that an expected utility argumentation like that from S. can solve these problems. A short summary of Kagan´s and my argumentation may be helpful.[24] We can start with D. Parfit, in whose famous scenario of the harmless torturer it is questionable whether a person P is morally right to be one of a thousand persons who, by pressing a button, make a very small contribution to increasing the electrical voltage to which another person Q is being subjected and which, when viewed overall, causes Q great pain.[25] But here the steps of each individual voltage increase should remain imperceptible to Q. P´s pressing of the button does not cause any altered perception of pain in comparison to the amount of measured voltage caused by my predecessor.[26] But in comparison with the initial and final states of all individual contributions combined, there is a great difference in the degree of suffering. How can the difference between 0 and 1000 be so grave when all intermediate steps are harmless? Kagan believes that such cases do not exist for a priori reasons.[27] But that means that you can solve problems like the Sorites paradox which is here implicated easily in one direction: "Vague terms like 'heap' or 'knowledge' … are semantically determinate so, in spite of appearances to the contrary, there is a sharp cut-off point to their application."[28] J. Nefsky offers another solution of the harmless torturers case according to which the small increase in voltage caused by P is not sufficient for leading to different perceptions of suffering. A minimal increase in voltage is simply not the right unit for causing perceptible suffering; just as in the Sorites problem, the grain of sand does not turn the heap into a hill that is the other direction of solving the paradox.[29] So there exist reasonable alternatives to Kagan's solution. If

---

[23] See Sinnott-Armstrong, op. cit., p. 304.

[24] See in more detail Kagan, op. cit.; see similar Gaverick Matheny, "Expected utility, contributory causation, and vegetarianism," *Journal of Applied Philosophy* 19 (2002): 293–297, Peter Singer, "Utilitarianism and vegetarianism," *Philosophy & Public Affairs* 9 (1980): 325–337, Derek Parfit, *Reasons and Persons (*New York: Oxford University Press, 1987) and others.

[25] See Parfit, op. cit., section 29.

[26] See Ibid., p. 80.

[27] See Kagan, op. cit., p. 131.

[28] Dominic Hyde (2018) "Sorites Paradox", *The Stanford Encyclopedia of Philosophy (Winter 2016 Edition)*, Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2016/entries/sorites-paradox/>.

[29] See Nefsky, op. cit., pp. 379-394. See Crispin Wright, "On the coherence of vague predicates," *Synthese* 30 (1975): 325–365.

Kagan's a priori point might not work, there remains an empirical *symmetry thesis*: in Kagan's argument exemplified by the purchase of a hen, the T^{th} purchase triggers the order (and thereby the production) of additional T products. For this he is relying on a market equilibrium concerning supply and demand:[30] A butcher "neither wants to fall behind demand nor end up with larger numbers of unsold rotting chickens."[31] But a market equilibrium is only justifiable where there are markets and even there, some counterexamples, exist, i.e. those caused by political influence on markets. The subsidies and governmental support measures related to the meat industry in the EU serve as a safety net that makes producers less vulnerable to fluctuations in demand. So I doubt that Kagan's argumentation is convincing.[32]

There may be validity to Kagan's thesis that, in cases governed by the market-place, the expected utility of an uncooperative action is negative as a rule.[33] But even if in Kagan's example, where there is a comparison of the expected utility of the behavioral alternatives as to whether or not to buy a hen from the butcher, the alternative "buy a hen" always has a comparatively negative expected utility: there are plausible reasons to doubt that this negative expected utility is *immense*.[34] Even if there is a probability that, for example, my specific "short" trip by car from Mannheim to Cologne could trigger a tipping point in the climate system, the high degree of damage that I could cause *is (almost) compensated for by the low likelihood that this event would actually occur*. There is apparently little reason to believe that the possible negative value of the utility expected from refusing to cooperate *would be immense*.[35] This is the reason why a vigorous dispute has arisen with regard to Kagan's question "Do I make a difference?"[36] Regardless of whether climate change and Kagan's et al. argumentation is a good example of a cooperative project with a low utility expected from individual contributions, cooperative projects of the harmless torturer type are just that, because individual contributions are by definition not perceptible there, and accordingly the effects of these cooperative activities are quite minimal when considered as individual contributions. Herein lies the fascination of the debate because, while individual actions are in themselves negligible, their accumulations give cause for concern.

---

[30] See Nefsky, op. cit., p. 369.

[31] Kagan, op. cit., p. 142.

[32] See Nefsky, op. cit., p. 369. See John R. Harris, Richard Galvin, "'Pass the cocoamone, please': Causal impotence, opportunistic vegetarianism and act-utilitarianism," *Ethics, Policy & Environment* 15 (2012): 368–383, p. 379.

[33] Understood as incremental difference, in comparison with the option of cooperation but not with an entirely different option. See Kagan, op. cit., p. 120.

[34] See Kagan, op. cit.

[35] See Mark B. Budolfson, "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response," *Philosophical Studies* 176 (2019): 1711–1724, p. 1718. See Douglas MacLean, "Climate Complicity and Individual Accountability," *The Monist* 102 (2019): 1–21, p. 6.

[36] I acknowledge that the more famous practical philosophers are among the defenders of this argument. But it could be the case that they are mistaken, whereupon consequentialism would be threatened in its entirety, as Kagan says in the next quotation. In that case, my argumentation would be insurance against relapse which demonstrates that, in this case as well, consequentialists doubtlessly have sensible options for argumentation and action.

The following problem arises for consequentialists in the case when individual actions do not make a causal contribution to the overall result:

"For if there are indeed cases that have this sort of structure (…) then consequentialism appears to fail even in its own favored terrain, where we are concerned with consequences (…). Intuitively (…) the acts in question need to be condemned because of the results that eventuate from everyone's performing them. (…) Yet despite this, it seems as though the consequentialist simply isn't in a position to condemn the relevant acts (…). The problem, in effect, is this: consequentialism condemns my act only when my act makes a difference. But in the kind of cases we are imagining, my act *makes* no difference."[37]

These problems regarding collective action can also to a significant extent be subsumed under this heading: decision amid uncertainty. I don't know whether my contribution will be useful in any way. What is interesting here is how consequentialists can deal with the fact that as an instrument for attaining knowledge, determinations of expected utility often leave us high and dry. We then have need of heuristics which complement the calculations of expected utility or we have to give up the project of consequentialism.

3 Great Event Heuristics and Ordinal Heuristics

What are we most likely to know about the utility of consequences when there is much we don't know, but if we can still in principle have knowledge about utility?

a) We can assess ordinal but not cardinal arrangements, at least when dimensions of extremely large magnitude (cf. the second and third premises of Hiller's calculation in section one) are involved. This is self-evident inasmuch as cardinal arrangements contain the same information as ordinal ones, but also further items: They indicate not only whether one entity is larger or smaller than another (which corresponds to an ordinal arrangement), but also the order of magnitude with regard to that difference. Thus cardinal estimates are more liable to error. With reference to Hiller's aforementioned utility analysis, we cannot measure the damage due to climate change in a cardinal numerical value (not even an estimated one), but can only evaluate this damage with respect to whether it is significantly larger or smaller than is the respectively offered comparative option A. I call the latter ordinal heuristics. Many analyses of expected utility, however, invite us to engage in what is most often an unsuccessful attempt not to assess utility values comparatively, but to estimate them per se.

b) We at least have knowledge about large differences in utility among options, i.e. about which options for action evince comparatively great or extremely little utility, because we can recognize large differences better than small ones. We know

---

[37] Kagan, op. cit., p. 107.

that utility is more enhanced by i) feeding a hungry person for a month than by paying 1000 German workers for the same time period a minimum wage of ii) €9.50 rather than iii) €9.49 per hour. The fact that we can assess comparatively large differences between consequences such as i) in comparison to ii) and iii) is the basis for our ethical and prudential decisions. Thus one can estimate with certainty whether in a soccer stadium A eight hundred (F1) or five thousand people (F2) are sitting (= dimension of visitors, namely "thousands," in which F2 is compared in terms of "hundreds" to F1), but it is not possible to estimate reliably whether in stadium B there are 780 (F3) or 881 (F4) persons (= exact difference between F3 and F4 in the same dimension of visitors, namely "several hundred").[38]

This means we can achieve good results with estimates of expected utility when there are comparatively large distances of utility between the options and we can then assess ordinal arrangements in a reliable manner[39]. Now it is a matter of expanding these insights into a heuristic. There are options for action whose comparatively large expected utility is fairly certain: for example, a contribution[40] of 30 euros to a good charity compared with the option of my cooperation or non-cooperation in boycotting a large oil company.[41] The sheer magnitude of the differences in utility achieved by a donation of 30 euros in comparison with an individual boycott that also costs me 30 euros makes a faulty assessment less probable. For 30 euros, I can restore sight to a blind person; but what effect is achieved by the 30 euros which I give, instead of to the boycotted company X, to company Y, because I now

---

[38] It seems superfluous to define an absolute standard for "big and little differences" here, because this may be *defined relationally* in terms of the "being very much larger or smaller" of an alternative: e.g. i) in the next-to-last example compared with ii) and iii).

[39] There is no great difference between estimates and analyses of expected utility per se, as can once again be demonstrated with reference to Hiller's calculation from section one, which works with estimates in all its premises. The primary difference will be between the precision of the results (cardinal or ordinal) and the expenditure of effort.

[40] Isn't the success of donations also dependent on cooperation? One could think that a contribution is all the more effective the more other contributions to the same goal are also received. The argument would be that the share spent on administrative costs declines with the volume of contributions, for example. But at the time of the donation, the donor to a good organization knows what the situation is with regard to administrative costs (Kagan cites an analogous scenario *in puncto* information: "New Haven Chicken Friends." See Kagan, op. cit., p. 127.) And sometimes small charities are more efficient than large ones. Someone for whom this degree of uncertainty is too large can, for example, consider contributing to foundations whose administrative costs, etc. are not covered by donations.

[41] See William MacAskill, *Doing Good Better – Effective Altruism and a radical new way to make a difference* (London: Random House, 2016), p. 178. There are several problems with charities certified by "Effective Altruism." There is reason to assume that American billionaires exhaust the capacity of these organizations to do good, so that additional contributions only lead to a reduction in the involvement of these billionaires, etc. See Mark B. Budolfson, Dean Spears, "The Hidden Zero Problem: Effective Altruism and Barriers to Marginal Impact," in H. Greaves, T. Pummer (eds.), *Effective Altruism* (Oxford: Oxford University Press, 2019), 184–201. Perhaps it would be advisable to avoid these certified top-charities in particular, but there are enough other, efficient charitable organizations, all of which have the advantage of not being dependent on collective action. Unfortunately, one has to gather information in order to be able to use this option of doing good.

buy somewhat more expensive gas there? The difference in utility between the two options catches the eye; the options a) "to contribute" and "not to boycott" and b) "to boycott" or "not to boycott" (and "not to contribute") lie in different dimensions of utility.

Here is an analogy situated somewhat differently and concerning not only purely quantitative estimates: We recognize with certainty whether something is blue or yellow, whereas we often cannot distinguish between shades of coloration such as light red and medium red (they are situated in a gray area that is liable to error). In order to avoid erroneous evaluations, we should avoid this gray area and concentrate on the paradigmatic cases which are certain (*paradigmatic move*). This is all the more true with regard to perceptions as well as to conceptual distinctions. There as well, we are better informed about the paradigmatic cases than about the marginal areas. What do we still call a heap, what is already a hill for us? On the other hand, the rising ground on which the "Wurmlinger Chapel" near Tübingen is built is clearly a hill. In our current context, paradigmatic cases are cases with a comparatively large[42] expected utility. Here the just-cited analogies end.

The smaller the differences in utility between the options, the greater is the liability to error on the part of our judgment, because our capacities for assessment and calculation (which are dependent upon estimates of subjective probabilities and utility values) quickly reach their limits. But also with the heuristic of great events being developed here, there is a need to evaluate whether the options are situated "close together" and accordingly within the "area of imprecision," or whether they evince great differences with regard to utility. An estimation of this, however, does not imply an exact evaluation of the options but only requires an "initial approximation," i.e. a fairly non-elaborate and inexact assessment of expected utility. This provides information about the utility dimension in which the options are situated.

Thus the heuristic worthy of recommendation would be the *great event heuristic* (GEH):

> GEH: Particularly with cases in which the dimension of the expected utility being targeted is small after a first approximation, one should not continue in involvement with the unclear options and in the complex endeavor to determine which option has which precise utility value; instead there should be a search for a third option with significantly higher expected utility, i.e. a great event (cf. the aforementioned example of contribution versus minimum wage). However, if there is no third option available at this time, check if you should turn to other problems, as it is not worth getting bogged down here.

If such a third option with higher expected utility exists, then the original options should be discarded, and there may very well be a need to examine whether the third option is optimal. Instead of considering whether a boycott or non-boycott makes a marginal difference with regard to welfare, I would be better advised, for example, to donate to a charity. Two options accordingly lose their relevance. A new question arises as to how the best course of action can be determined. If necessary, there can

---

[42] Or small expected utilities; but that is irrelevant for the context.

be a more precise analysis of costs versus utility here. But a warning must be issued about the mistake of continuing to focus on false options and attempting to choose the correct action through more precise analyses of the differences within a narrow range of utility dimensions.[43]

One could make a critical comment in general by arguing that what matters is not the size of the differences but the number of uncertainty factors in order to make reliable predictions, for example. If I know for certain how the future will look, I can base decisions also on minimal differences.[44] But my basic thesis is that, with regard to small differences, it is not known with certainty whether they exist and what their exact magnitude is. Uncertainty and size are not inherently connected, but they are linked by the laws of estimates and calculations: with both, *uncertainty correlates positively with a minimal magnitude of the utility difference*. Great event heuristics reacts to an epistemic problem (see below: the thesis of imprecision), and of course it is senseless to eliminate an epistemic problem by means of a theoretical decree. As a rule, I don't know what effect small differences have, and I must make a huge, complicated effort to achieve sufficient certainty in this regard; even this may be in vain. The great event heuristic seeks to relieve me of this obligation by arguing that in cases of minor differences (in the gray zone), there is nothing to be gained from the (perhaps not at all successful) analytic effort, and that an attempt should be made to depart from the gray zone.

Two arguments involving great event heuristics[45] offer a succinct formulation of the heuristic under discussion here:

GEHA (Great Event Heuristics Argument) 1:

1. Precise estimates and analyses should (and sometimes could) not be used as methods of determining the expected utility of options in which there is a large susceptibility to error on the part of these methods[46].
2. In cases in which, after a first approximation, two possible courses of action are distinguished solely by a small unit of utility (i.e. they are situ-

---

[43] The situation could remain problematic if all options lie within the gray area, but then it could be said: "Don't concern yourself with these options but wait until a clearly better one becomes available." This could be a situation, however, in which I am no longer able to comparatively determine the magnitude of the utility dimension. An attempt could be made to introduce absolute standards in such situations: the "magnitude" is best measured in terms of the average expected utility which can be realized by the action of a "normal" individual (i.e. neither the President of the USA nor a homeless person, but instead an average middle-class citizen, for example).

[44] J. Lenman's example is the invasion of Normandy during the Second World War. If the section of beach most promising of success for the landing had not been known and it had only been known that, if beach A were chosen, the leg of a dog in Amsterdam named Spot would not be broken (which would otherwise happen), then the section of coastline could only be selected on the basis of the comparatively minimal information concerning Spot's leg. Lenman argues that all other consequences can be assumed to be equally probable and accordingly irrelevant with regard to the decision to be made. See James Lenman, "Consequentialism and Cluelessness," *Philosophy & Public Affairs* 29 (2000): 342–370.

[45] This term is defined by the following conclusions C1 and C2.

[46] Analyses as well, because they also work with subjective probabilities. Sometimes differences are so small that they cannot be detected by any method. In these cases, there can be an unavoidable vagueness.

ated in a gray area), the susceptibility to error on the part of these methods for determining expected utility is de facto large.

C1: Precise estimates and analyses of expected utility should (and sometimes could) not be used for determining the expected utility values of options which lie in the gray area and in which, after a first approximation, two options for action are distinguished solely by a minimal utility unit (= thesis of imprecision).

GEHA 2:

1. Precise estimates and analyses should not be used as methods of determining the expected utility of options which lie in the gray area.
2. Instead one should use as a utility maximizer expected utility values which are as large as possible; in other words, a rough estimate (i.e., in a first approximation) should be undertaken with regard to options which have significantly higher expected utility than the options in the gray zone.
3. DEF: Great events are defined as a set S_1 of paradigmatic options for action characterized by the fact that, after an initial approximation, the expected utility of the elements of S_1 is significantly larger than that of the elements of the set of options for action situated in the gray zone.
   C2: As a utility maximizer, one should attempt a rough estimate of great events (paradigmatic move).

The picture is complicated by the fact that between various great events, the thesis of imprecision often dominates and gives rise to a new inestimability which may be called gray area 2. For example, is the utility of a donation of 30 euros to a good organization D(O) larger than that of a contribution of 30 euros to another good organization D(P)? Both options are clearly better than the boycotting of an oil company by a single person (i.e., an option in gray area 1, and S(O) and S(P) are comparatively great events because they clearly leave gray area 1 behind), but both lie close together with regard to their high utility. They clearly protrude from gray area 1, but which is better? In order to emerge from this new gray area 2, great event heuristics does not help further, because there exists no third option with clearly larger utility values. An expected utility analysis could help in a case where the anticipated difference of the utility values justifies the expenditure of effort. In the case of ongoing imprecision, a thesis of indifference could be valid: both contributions are justified if neither of the options clearly performs better. Thus it makes particular sense to use the great event heuristics when there is an overlooking of large potentials for utility (for example, from contributions) which could be achieved instead of wasting time with false alternatives of minimal utility (e.g. cooperative efforts).

One may be inclined to label heuristics of this sort as trivial, because the utility maximizer has always given consideration to them. If one seeks to maximize utility, it seems appropriate to first seek out and work through great events rather than to take a path via the summation of small events. (These include at least the aforementioned collaborative projects of the "harmless torturer" type.) But a look at actual practice has something different to tell. There is the counter-position, for example

in climate ethics: "My only point is that it is easy to underestimate the marginal effects of individual actions, which turn out to be non-negligible."[47] Here philosophers argue over the damage caused by a single drive in a SUV or one solitary act of consumption, engaging in calculations such as those of Hiller and others cited at the beginning.[48] Is this hair-splitting or a solvable arithmetic problem? Hiller suggests the latter: "Going for a Sunday drive has the expected effect of ruining someone's afternoon. To use non-technical moral terminology, that is not a very nice thing to do."[49]

These considerations are not only the starting signals for a theoretical investigation, but also for the formulation of imperatives of the type "Buy products of the type X!" and prohibitions such as "Avoid emissions actions of the type Y!"[50] These may refer to quite minimal amounts of expected utility in collective action cases; in this regard, the thesis is here being put forward that in these debates, sight has perhaps been lost of great events. In such situations, there is a need for effective heuristics. As cited earlier, Feldman has described the standard path of expected utility determinations, wherein the possibility of utilizing heuristics is not even mentioned, even not in his proposed solution for the concomitant problems.[51] I am of the opinion, however, that they are of primary importance and should preferably be used.

4 Big versus Great Events

T. Cowen expresses at a first glance similar thoughts with his "big event consequentialism."[52] He as well proceeds from the idea of imprecision: "The Principle of Roughness (…) leads us to discriminate against relatively small benefits and losses. (…) Large upfront benefits are less likely to be overwhelmed by the roughness of our comparisons, and thus it is compelling to pursue large benefits."[53]

But at which point does Cohen's approach differ from mine, and what can be learned from this difference? Cowen considers large benefits or big events to always be so huge that the actions of a normal individual are utterly incapable of influencing them. Cowen conveys this impression by frequently speaking of big events as if all that mattered were only the magnitude of the project and not the individual expected utility. He speaks about "fighting smallpox," "fighting hunger," "defending against

---

[47] Hiller, op. cit., p. 20.

[48] See John Nolt, "How Harmful Are the Average American's Greenhouse Gas Emissions?" *Ethics, Policy and Environment* 14 (2011): 3–10, p. 3.

[49] Hiller, op. cit., p. 20.

[50] Lumer, op. cit., p. 104. See Christopher Morgan-Knapp, Charles Goodman, "Consequentialism, Climate Harm and Individual Obligations," *Ethical Theory and Moral Practice* 18 (2015): 177–190, p. 190.

[51] See Fred Feldman, "True and Useful: On the Structure of a Two Level Normative Theory," *Utilitas* 24 (2012): 151–171.

[52] Tyler Cowen, "The Epistemic Problem Does Not Refute Consequentialism," *Utilitas* 18 (2006): 383–399.

[53] Ibid., p. 390.

terrorists" and so forth.[54] In the case of important contributions to these projects, one is indeed more justified in assuming a great utility than when deciding to go to the movies this evening. But that could imply that utility can only be clearly attested there where the individual normally has no influence. Politics has an immense scale of influence; ordinary individual behavior almost never evinces such great differences of utility. And there can be marginal contributions to big events: *The type of event in and of itself does not determine the size of the individual involvement.*

In order not to end up in such implausibilities, I speak of great rather than big events, which entails *a systematic shift towards the magnitude of expected utility from an individual perspective.* As a rule, it is of course clearly destructive when a person commits murder, and useful when an individual helps others. In any case, the "size" of the expected utility does possess a certain relevance on the individual level. If each month I donate 30 euros against blindness, I will be acting better than if I add up a succession of projects whose utility is uncertain and, if present at all, extremely small: for example, in that I trim the hedge in front of my house because I hope that passers-by will enjoy the cared-for shrubbery. Hence I formulate the individual great event principle IGEP as follows: "It is a matter of identifying for the respective subject or the respective institution the largest attainable expected utility and of concentrating on it."[55]

One could be of the opinion *that this is self-evident for utility maximizers.* Already above with the example of climate ethics, it was demonstrated that this is de facto not the case. Nor are heuristics present in the lives of many ecologically-minded consequentialists when they take great pains to reduce emissions and to become vegetarians.

5 Collective Action Problems and Great Events

So let it be summarized with complete clarity:

1. Great event heuristics teaches that every protagonist must show self-restraint there where, on the basis of a rough estimate, it is not the case that he can realize great events.[56]
2. No great event can be discerned with regard to a participation in social collaborations in which the scanty probability of one's own cooperative contribution having an actual impact (almost) outweighs its potential utility.
3. DEF: I call social collaborations such as these "symbolic collaborations."
   C: Every protagonist should show self-restraint with regard to symbolic collaborations.

---

[54] Ibid., pp. 387, 394.
[55] It is not fundamentally important whether Cowen means this or only formulates it infelicitously.
[56] To an exception in the next section.

But is it not the case that many great successes came about through cooperative projects? Doesn't conscious consumption by purchasers alter the behavior of companies? Wasn't it once already possible through a boycott to force Shell to break down and depollute the drilling platform "Brent Spar" instead of simply sinking it into the ocean? But should I participate with resources which I only have available once in a large—and if so, only minimally capable of being influenced—project, or should I select an undertaking with a greater benefit, for example through donations? The latter should be the case, because everyone is responsible for making as effective a use as possible of his always limited moral commitment. Hence one should refrain from symbolic collaborations. It is true that this endangers the success of such projects. If many people were to act in this way, that would be fatal for the project. But most of the time, others act independently of me; as a rule, they take no notice of my decision. Of course, it is even more dangerous to publicize this thought; because in this way, it may be that such a number of subjects is influenced that together they are in fact capable of having an impact. But even if this is bad for the collaborative project, *it can be better overall* if with their courses of action, all protagonists make big differences instead of at best small ones. Moreover, a successful collaboration is not the only way to a goal of cooperation which, in the case of absent participation by individuals, can be compelled by administrative decree. Therefore: if necessary, *it is preferable to be a free rider* and to avoid symbolically applying the scanty resources of moral action, and to direct one's efforts towards greater effects. A strict deontological prohibition of free riding would not be a feasible alternative, because cases can quickly be found in which an act of unconditional cooperation would demand so great a sacrifice that it would in any case not be insisted upon if one believes that morality exists in order to help people.

Of course you may change the situation if you *engage as a voter*. Let us therefore take the example of climate change and the individual's behavioral response again. W. Sinnott-Armstrong writes: "It is better to enjoy your Sunday driving (in a SUV, B.G.) while working to change the law than to make it illegal for you to enjoy your Sunday driving."[57] He argues that individual action is not a necessary or sufficient precondition for causing climate change, and we would be better advised to pass laws against climate violations than to avoid emissions. As said before, this argument is countered by Hiller and others who support Kagan.[58] At a first glance, great event heuristics rejects *both* these positions, if we translate "make illegal" with "elect representatives who enact certain laws." With respect to neither emissions nor elections do I (in a normal situation) play a determining role. An analysis of elections shows that as an individual voter, I seldom make a difference; it may very well be that here I am chasing after minimal variations in utility.[59] W. MacAskill and others show that it often is a matter of the context. Voting can have a large expected utility, but this is frequently lacking.[60] (In the next section we come back to voting under the aspect of low-cost-strategies.)

---

[57] Sinnott-Armstrong, op. cit., p. 304.

[58] Hiller, op. cit. Morgan-Knapp/Goodman, op. cit.

[59] See Anthony Downs, *An Economic Theory of Democracy* (New York: Harper & Brothers, 1957).

[60] See MacAskill, op. cit., Chapter 6.

But how can I as an individual contribute to extensive projects? The greatest effect of individual actions is often considered to lie in donations to certified charities, to the extent that one has any financial means available for improving the world.[61] But one has to gather information as we learn from Budolfson and Spears.[62] A single donation that gets through to the recipient has the advantage that it can greatly enhance its utility for him (the alternatives of vision or blindness), and this is not the case with symbolic collaborations: "Whether our concern is the global poor, climate change, or animal welfare, we´ve seen that the decision about how much and where to donate is much greater, in terms of impact, than the decision about which products to buy."[63]

Additionally a thesis of imprecision similar to the one that I established above can simultaneously free utilitarianism from normative objections such as the argument that an extremely minimal utility provides the justification for favoring e.g. a sadist: "If the calculation gives rise to only one small utility unit more for the sadist, one must allow torture." Judgments such as these are often applied against act-utilitarianism.[64] They are perhaps conceivable if cases are sufficiently isolated and shortened into manageable sets of consequences. But they are theoretically deficient, because they disregard an inviolable region of imprecision which can be present in real-world cases. We are scarcely and sometimes not at all capable of knowing whether a small difference in utility exists between two events; one should accordingly refrain then from making a judgment with regard to the difference in magnitude. The size of this zone of imprecision, just as with all zones in similar cases (in the case of both perceptions and concepts), is unfortunately also itself indefinite, or only capable of being investigated in anthropological and psychological terms.[65] In cases as this one, in which the consequences of my action are undetectably small, it may be impossible to decide which of my options has the best consequences. This impossibility to decide can exist both at the level of decision procedure and at the level of rightness. If not difference can be detected even in principle, there is not basis for determining the rightness of actions.

So to summarize: Has the problem of cooperation been solved? Is consequentialism not in fact simply recommending a change of perspective? As fascinating and important from the perspective of the philosopher the harmless torturer problem of whether I make myself guilty or not when I press the button is, it is a mistake

---

61 See Ibid., p. 167.

62 See Budolfson/Spears, op. cit.

63 And even if Kagan's argumentation were correct, that doesn't change the psychology. Because of the complexity of the interconnections, many consumers will continue to *fear* that their contribution to a problem as huge as climate protection will perhaps have no effect. Motivation is higher when it may be expected with relative certainty that the individual contribution will change something, at least for a concrete person. MacAskill, op. cit., p. 178.

64 See Brad Hooker, *Ideal Code, Real World (*Oxford: Oxford University Press, 2000), p. 156.

65 See Cowen, op. cit., p. 394.

from the perspective of great event heuristics to remain focused on it.[66] With such minimal differences beyond any possibility of distinction by perception as in Parfit´s example, it is idle to decide that one should act so or so, because such decisions are rendered irrelevant by the thesis of imprecision from GEHA 1. This is an unromantic solution to a problem which legitimately catches the interest of the tinkerer, but not of the utility maximizer.

But is that all that needs to be said here? Must we in many cases cancel projects such as consumer ethics, the individual's change in his emission behavior and his political participation in the form of voting because they tie up energies which would be better used for great events, e.g. donations?

6 Collective Actions and Low Cost Tolerance

There are arguments in favor of not completely abandoning the aforementioned strategy of cooperation, *but instead of continuing it there where the costs are situated in the area of imprecision*. Why is that so? Should we not, when we have available a means such as deliberate individual donations which is superior on the basis of utility estimates, make an energetic commitment to that means? Isn't that a maximization of utility?

Astoundingly, the answer is "No." 1) The arguments of the first section showed us our susceptibility to error. It could be that we also make mistakes when using great event heuristics (e.g. with the "first approximation"). Then we cause possible irreparable damage to cooperative projects whose goal is very important. So we should keep a *back door* open if this happens to be possible.

2) It is possible if the differences in costs (of buying x and buying y) lie comparatively in the dimension of negligibility. Regardless of how the extent of benefits and the probability of their occurrence appear: If they cost us nearly "nothing" extra, then they don´t matter. The cost factor plays a special role in the determination of expected utility. As long as the cost differences in this regard are worthy of mention, I should not change my emissions behavior as long as the overall expected utility does not turn out to be larger. But I can undertake "at nearly no cost" any number of measures whose expected utility is not particularly high. This can be applied at least to political voting. Normally in industrial nations, this only costs us the time to mark an "X" and mail it in a letter. These costs are comparatively small, whereas the utility could be immense. A. Maltais writes: "However, voting can be morally obligatory because its very small impact is a corollary of its very small cost to the voter."[67]

---

[66] Simply enlarging the achieved differences in utility does not save the example from criticism, because this has the effect of undermining the fact that it is not possible to distinguish between the adjacent steps, and thereby of altering the type of collective action problem being addressed.

[67] Aaron Maltais, "Radically Non-Ideal Climate Politics and the Obligation to at Least Vote Green," *Environmental Values* 22 (2013): 589–608, p. 602. As for information costs: they can be high but sometimes the alternative is clear as could be. For example, I do not need to first inform myself about the exact programs of right-wing radicals. Perhaps I should simply participate in elections in which I am sufficiently informed and the results could be extremely close. See MacAskill, op. cit., chapter 6.

In this way, the aforementioned quotation from Sinncott-Armstrong, in which he calls for political involvement, can be corroborated, at least with regard to that part of the political demands which requires participation in elections.[68] Particularly if we are skeptical with regard to precise analyses of costs and utility, we should here not excessively pursue a calculation of cost versus utility relationships in the case of marginal costs with great utility potentials. And we should be particularly skeptical with regard to "precise" analyses of marginal differences, because, as the thesis of imprecision teaches us, we are scarcely able to recognize small differences in utility, so that a certain generosity recommends itself here: *I call it low-cost tolerance.* This opens the aforementioned back door.

Participation in collective action projects can be pursued as an alternative causal strategy in addition to great events (e.g. donations) when they lie in the low-cost tolerance region: for example, voting. Moreover, the success of some collaborative projects provides *causal contributions* to the success of other strategies: for example, donations. The degrees of utility that can already be attained in the short term (e.g. through donations for fighting the poverty of farmers in the rain forest so that they no longer have to cut down trees, which has a positive effect on reducing poverty, maintaining species and protecting the climate) become more successful or acquire further *multiple effects* when the long-term developments (e.g. political climate protection achieved by elections, and ensuing agreements for preserving the rain forest in Brazil) likewise proceed in the correct direction. Cooperative efforts and donations can enhance each other.[69] The expected utility of the collaboration must no longer be evaluated solely on the basis of its direct consequences; attention should likewise be given to multiple impacts upon other measures (e.g. donations). This is self-evident in the theory of the utility maximizer, but the fact is that it is often neglected. The expected utility of cooperative efforts *could* accordingly be higher than it seems at a first look, and the cost differences of low-cost cooperation would remain low. Further heuristics other than the two which have been defended in this essay are conceivable: for example, giving reversible steps priority over irreversible ones, so as to minimize certain risks.

7 Summary: Cooperation and Heuristics

In conclusion here are a few heuristic rules for decisions which are certainly not complete:

1) When possible, use unproblematic principles of the first level in the framework of Hare's two-level morality.

---

[68] Political activism or involvement (e.g. Greta Thunberg) has other costs, and its utility potentials would require a more detailed analysis than is possible here.

[69] Of course we must change our live-style in the long run. But the recommended strategy allows us to do this later and to give technology opportunities to save things like our beloved individual mobility. So motivation to help is conserved because making a donation is doubtlessly less deterring and more motivating that radical changes in behavior. To demand to do all things at once might be over-demanding and motivation-killing.

2) If that is not possible, search on the critical level for great events which can be identified with rough estimates of expected utility.

3) If it appears beneficial, try to analyze the expected utility of the great events in more detail.

4) Not least of all, in recognition of the individual's own fallibility, also employ strategies which are less favorably positioned in the estimate of expected utility, if they are low-cost options.

5) Also recognize that the issue is frequently overlooked as to whether measures can lead to multiple effects and synergies.