

## Individual Climate Risks at the Bounds of Rationality

Avram Hiller  
[ahiller@pdx.edu](mailto:ahiller@pdx.edu)

Forthcoming in *Risk and Responsibility in Context*, Adriana Placani and Stearns Broadhead, eds., Routledge (2023). Please cite published version when available.

*Dès qu'il y a vie, il y a danger.*

Madame de Staël (1845, 564)

### 1. Introduction

All ordinary decisions involve *some* risk. If I go outside for a walk, I may trip and injure myself. But if I don't go for a walk, I slightly increase my chances of cardiovascular disease. Typically, we disregard most small risks. When, for practical purposes, is it appropriate for one to ignore risk? This issue looms large because many activities performed by those in wealthy societies, such as driving a car, in some way risk contributing to climate harms. Are these activities morally appropriate?

In this paper, I will argue that it is appropriate to ignore many small risks. I am not the first to argue for this conclusion. However, the reasons that I give for ignoring small risks differ to some extent from those identified by others in some recent debates. In particular, I will argue that because our rationality is *bounded*, it is impossible for us to include every small risk in our decision-making process, and so we may reasonably use heuristics to guide many decisions. Although our use of heuristics allows for the reasonable ignoring of some risks, and perhaps explains why one might be inclined to think that individual climate-related risks are negligible, the

main aim of this paper is to show that even reasonable use of heuristics does not permit the *general* ignoring of climate-change-related risk by individuals.

The other main aim of this paper is expository. Philosophers have engaged with issues related to the present one in great detail, especially since Derek Parfit's *Reasons and Persons* (1984).

Although in this paper I advance a particular thesis about when it is appropriate for individuals to ignore their greenhouse gas (hereafter GHG) emissions' climate-related risks, the relevant literature is vast and multidisciplinary (including climate science, environmental economics, behavioral economics/psychology, decision theory, ethics, and metaphysics). Because the evidence required to assess individual climate-relevant obligations is so large and disparate, I have attempted to bring together several literatures that pertain to the question. Although I am not able to be comprehensive in my discussion of the relevant literatures, I hope that readers will benefit from my efforts to bring together considerations from several different fields even if they disagree with my conclusions.

I should note that economists and decision theorists sometimes distinguish between *risk* and *uncertainty*. The former is when there are specifiable probabilities of an outcome occurring, and the latter – *Knightian uncertainty*, (from Knight 1921) – is when there are not. In this paper, I will not distinguish between the two, and will use the notion of *risk* to cover cases when there is some possibility (whether it be precisely known, imprecise, or unspecifiable) of some negatively-valenced outcome occurring. Cases of individual climate risk are ones of Knightian uncertainty.

## **2. "It Makes No Difference," *again*?**

A number of philosophers (Sinnott-Armstrong 2005, and later Kingston and Sinnott-Armstrong 2018, as well as Cripps 2013; Nefsky 2012; Sandberg 2011; Gesang 2017) argue that, despite the real existence of global climate change and its harms, individual actions do not make a difference in regard to climate change. The general idea, in Kingston and Sinnott-Armstrong (2018), is that the effects of one individual emitting greenhouse gases are so small and diffuse in the global causal structure of climate change that it makes no difference whether one individual does or does not emit greenhouse gases at that scale.<sup>1</sup>

In this paper, I will not engage directly with these arguments to any significant extent. My own view, mostly in line with Broome (2019 and 2021; also see Hiller 2011a and 2011b; Morgan-Knapp and Goodman 2015; Nye 2021), is that (A) GHG-emitting actions should be held to have some non-negligible amount of *expected* negative disvalue and that (B) this expected negative disvalue is morally relevant. Roughly speaking, Broome (2019) shows that individual GHG-emitting activities have a strictly increasing *expected* harm function even if they do not in fact cause harm or trigger a harm threshold. This is not to say that actions that emit GHGs are never all-things-considered permissible; to determine whether they are permissible depends upon one's overall normative views and other relevant facts. The maximizing consequentialist, for instance, will hold that a GHG-emitting action is permissible if it is the best action one might take. Perhaps the expected *benefits* from certain GHG-emitting actions are high, or perhaps one has no other *better* options than to emit GHGs. But the debate surrounding individual climate ethics following

---

<sup>1</sup> This debate echoes one in animal ethics, where some philosophers argue that we do not have an obligation to reduce our meat consumption despite the fact that animals are deserving of moral consideration; see e.g., Fischer (2019, Ch. 4) and Nath (2021).

Sinnott-Armstrong (2005) has largely been occupied with the question of whether there should be *anything* on the negative side of the ledger from the expected climate harms due to individual actions, and it is this question that I take to be answered in the affirmative.

Although I will raise some concerns about some details of Broome's analysis, in this paper I will primarily take a broader look at the debate. Is standard expected value theory applicable to individual climate ethics in the first place? At the very least, the application of expected value theory to climate change is *fraught*: in one formulation, it involves taking a tiny fraction – the proportion of the total amount of global GHG emissions that one individual is responsible for – and multiplying it by a huge value – the total amount of harm that can be expected to occur (in the form of a harm function) given various climate scenarios. While Nolt (2011), Broome (2019), and Hiller (2011a and 2011b) unapologetically perform this multiplication, and hold that the resultant product is a non-negligible value, the very consideration of a tiny risk of harm may be objectionable for independent reasons.

Within the field of risk analysis, for example, some have argued that certain small risks, referred to as *de minimis* risks, may be reasonably ignored. Peterson (2002) gives the first extended philosophical analysis and critique of *de minimis* risk (also see Adler 2007 and Lundgren and Stefánsson 2020 for related criticisms). The idea behind *de minimis* risk is that some risks are too small to merit consideration by the law. For instance, the law need not require that all carcinogenic items be omitted from food items, because potentially all items may raise cancer risk by the tiniest of amounts. As Peterson notes (2002, 48), initially the notion of *de minimis* risk was intended as a response to an extreme form of the *precautionary principle*, since a *de minimis* risk

principle says that precautions need not be undertaken when risks are extremely small. Still, the notion of *de minimis risk* can also be used in the context of a less extreme form of the precautionary principle, because when there is a tiny risk of a potentially large catastrophe, policymakers may still wish to ignore that possibility. Even Cass Sunstein, perhaps the most ardent advocate of cost-benefit analysis, supports the employment of a *de minimis* risk principle (2002, 193-195; 214-216).

As Peterson argues, the notion of a *de minimis* risk admits vagueness – there is no sharp boundary between the *de minimis* cases and the non-*de minimis* cases. Lundgren and Stefánsson (2020, 913) further argue that in cases where there is a tiny risk to one option, and no benefit, then it would be wrong to impose the risk. These arguments seem to me to be correct. The literature on *de minimis* risk has largely pertained to risks from governmental policies, and I will bracket it for the moment.

A related literature has arisen recently regarding the notion of *discounting* small probabilities.<sup>2</sup> Kosonen (2021), Monton (2019), and Smith (2014) have argued that we may rationally discount tiny probabilities. For instance, Bostrom's (2009) case of Pascal's Mugging seems to show that without discounting small probabilities, expected value theory is susceptible to a seemingly absurd result of indicating that it is best to sacrifice finite goods for an astronomically unlikely promise of enormous rewards. Balfour (2021) explicitly connects this argument to existential risk for humans, arguing that we must thus make extreme and ridiculous efforts to reduce existential risk, and suggests that expected value theory should be abandoned for this reason.

---

<sup>2</sup> It should be noted that this notion of discounting differs from *temporal discounting*, a common notion within climate ethics.

In response, a number of commentators have responded that discounting itself leads to absurdity. For instance, both Ebert et al. (2020, 438-439) and Barrington (ms.) argue that partitioning the outcome space of a choice into multiple partitions may turn an intuitively wrong choice into an appropriate one if discounting of tiny risks occurs. If two possible bad outcomes each independently have tiny risks, then these can be ignored according to a discounting principle, but if the outcome were simply a single outcome with the sum of these risks, then it could not be properly ignored. But this ramification of discounting seems inappropriate.

Gesang (2021) argues that individual actions are subject to inefficacy and inscrutability concerns, and thus are not subject to expected value analysis. At the same time, according to Gesang, *large* ones *do* make a significant expected difference. However, for reasons related to those just discussed, this argumentative move must be mistaken. For instance, Hiller (2011, 355) notes that drives themselves can be partitioned into smaller concatenated actions. If there is a threshold below which risks are morally negligible, then one could avoid having responsibility for many of one's culpable doings simply by dividing them into smaller ones. If going on a drive across the United States is above the culpability threshold for climate-based emissions risk, but driving shorter distances is not, one could simply plan to leave from one coast, drive from town to nearby town, eventually happily arriving at the other coast without the burden of a guilty conscience. But this seems absurd. Instead, traditional expected-value theory seems vindicated, because it simply adds up the disvalue of the smaller drives into the same sum as the single longer one. *Causes* in group or large event phenomena, on the other hand, can't be decomposed into equal partitions. What my particular emissions will in fact trigger may be different from what your particular

emissions will trigger. But *expected value* of large actions can be divided into partitions when there are no known reasons for making an exception for a particular marginal contribution.

Additionally, Barrington argues (ms., §3), similarly to Lundgren and Stefánsson, that when there are no *other* relevant factors, it would be wrong to ignore a tiny probability of harm simply because it is small. Indeed, it does not seem right that the mere *minimal* nature of some risks could, entirely on its own, be sufficient reason to ignore them.

Another longstanding argument in the literature has been more broadly viewed as successfully undermining traditional expected value theory – the *Small Improvements Argument*. Although the relationship between the small improvements argument to individual climate ethics might seem at first glance distant,<sup>3</sup> what I will argue in the next two sections is that lessons from the psychology of decision-making that are revealed by a close look at the small improvements argument can help shed light on individual climate ethics.

### **3. The Small Improvements Argument**

It is a natural reaction in certain cases to think that certain small differences in outcomes should make no difference in the choiceworthiness of options. Joseph Raz (1986), following Ronald de Sousa (1974), gives one such case. Imagine a situation where a person is deciding upon a career as a philosopher or a career as a lawyer. Assume that neither option is better than the other with

---

<sup>3</sup> Broome (2019, 124) has some discussion of the Small Improvement argument, but Broome's use of the SIA is quite different from the one in this paper.

respect to success or desirability of career for the person, and the person finds it extremely difficult to choose. Plausibly enough, argues Raz, if the person learns that the legal career has a very slightly better salary than the person had previously considered, it is *still* the case that neither career choice is better than the other for the individual.

Several philosophers have concluded from this example that values are *incommensurable*, a claim that deserves fuller treatment elsewhere. What is relevant here is that this example purportedly shows that foundational principles from classical expected value theory do not hold. In particular, here are two core principles: (Let  $V(X)$  mean the value of state of affairs  $X$ .)

*Completeness*: for any states  $A, B$ , exactly one of the following is the case:

$$V(A) = V(B); V(A) > V(B); \text{ or } V(A) < V(B)$$

*Transitivity*: If  $V(A) = V(B)$ , and  $V(B) = V(C)$ , then  $V(A) = V(C)$

In Raz's case, let  $P$  denote the individual's career as a philosopher; let  $L$  denote the individual's career as a lawyer; Let  $L+$  denote a career as a lawyer, but with it being slightly more lucrative than in  $L$ . The intuitive set of claims are that

$$(1) V(L) = V(P)$$

$$(2) V(L+) = V(P)$$

$$(3) V(L) < V(L+)$$

but by (1) and (2) and transitivity,

$$(4) V(L) = V(L+), \text{ violating completeness}$$

Instead, Ruth Chang (2002) argues that the proper characterization of the situation is to say that the values of  $L$  and  $P$  are *on a par* (see Andreou 2015 for a more recent defense of the notion of parity).



Interestingly, one way to initially understand the Kingston/Sinnott-Armstrong view is to claim that if  $D$  is a typical Sunday drive, and  $D+$  is a Sunday drive in a vehicle that emits no GHGs, that  $V(D) = V(D+)$ . But, intuitively at least,  $V(D) < V(D+)$ , perhaps violating completeness. Although one shouldn't say that  $D$  and  $D+$  are *on a par*, Kingston and Sinnott-Armstrong would likely hold that for any choice when one is faced with a decision between  $D$  and some *other* option  $O$ , whether to choose  $D$  or  $O$  need not involve consideration of  $D$ 's GHG emissions, and thus the relevant features of the individual's deliberation should be the same as in a deliberation between  $D+$  and  $O$ . The choiceworthiness of going for a drive relative to some other option should be the same whether or not the drive emits GHGs.

One early response to the Small Improvements Argument is that it is an *epistemic* issue (Regan 1989, 1059-1061). Regan's idea is that even if we intuitively think that both  $V(L) = V(P)$  and  $V(L+) = V(P)$ , they are not both true, and this is because we are not properly grasping the relevant fine-grained states of affairs in question. We are simply uncertain of  $V(L)$ ,  $V(L+)$ , and  $V(P)$ , and are not really evaluating them.

I should mention that Ruth Chang responds (2002, 669-670) to the kind of concern raised by Regan by noting that the small improvement phenomenon occurs not just for major life decisions but also small decisions, like ones in which one decides whether to have a cup of coffee or tea. If one is undecided between the two, and then one hears that a slightly better tea is available, that will not necessarily sway one to choose the newly available tea over coffee. It is not my aim here to delve into all the details of the Small Improvements Argument, but I should note that this example

has never seemed convincing to me. When one evaluates coffee against tea, it is not the coffee and tea that are intrinsically valuable; what matters are the experiences that they will produce in the drinker. But it is unclear how carefully the drinker can in fact anticipate the full set of experiences that they will have upon drinking coffee or tea, especially given that the value of aesthetically pleasing experiences is modulated by the context one is in. Even a choice between familiar coffee and familiar tea can have significant uncertainties in it in any new circumstance, and so the epistemic move still seems appropriate.

The general approach that I'd like to take in responding to the Small Improvements Argument is similar to that of Regan and of Anderson (2015). But I will express it using some notions from the field of behavioral economics.

#### **4. Bounded rationality and minimal risks**

##### ***4.1. Bounded rationality***

The concept of *bounded rationality* is familiar to many (see Gigerenzer 2021 for a helpful history). Herbert Simon (1955), and later Amos Tversky and Daniel Kahneman (see 2011) and many others argue that there are significant constraints on human abilities to reason. Some of these constraints can primarily be seen as endogenous to the human mind: humans can sometimes be slow in arriving at answers, and we are susceptible to biases that cause us, in certain contexts at least, to regularly provide incorrect answers. It is, famously, a controversial issue whether this should lead us to think of humans as being to a significant extent irrational or, rather, to understand rationality as necessarily contextual/ecological (as in Gigerenzer 2000). Other constraints are best

seen as exogenous: we often have limited time and evidential resources to make judgments and decisions. Here, I will remain neutral on the question of the rationality of endogenous constraints on human judgment and decision-making and instead focus on constraints that are primarily exogenous.

Kahneman and Tversky, as well as Gigerenzer (2000), emphasize that people employ *heuristics* in making judgments and decisions. Although oftentimes heuristics are commonly thought of as *rules of thumb*, and Gigerenzer (2000) understands heuristics as tools, I'd like to employ a definition given by Kahneman in his later collaborative work with Shane Frederick (2002). As Kahneman (2003, 466) summarizes, "A judgment is said to be mediated by a heuristic when the individual assesses a specified *target attribute* of a judgment object by substituting a related *heuristic attribute* that comes more readily to mind." So rather than use cognitive resources in assessing a complex state, individuals use heuristic attributes, which substitute for the more complex state and are easier to assess.

#### ***4.2. A response to the Small Improvements Argument from bounded rationality***

It is not hard to see how heuristics can be employed in an epistemic solution to the Small Improvements Argument. "Life as a Lawyer" and "Life as a Philosopher" are heuristic attributes, which stand in for more complex states. When comparing  $L$  to  $P$ , one compares these two heuristic attributes, and deems them the same in terms of their value. When comparing  $L+$  to  $P$ , one still compares these *exact same* two heuristic attributes. Even if a fully specified fine-grained state of

affairs of  $L+$  is better than a fully specified fine-grained state of affairs  $L$ , we are not comparing either of those to  $P$ .<sup>4</sup>

Perhaps one reason why no one has understood the Small Improvements Argument in these terms is that Kahneman and Tversky and Gigerenzer are all explicit that heuristics are *fast and frugal*. But deliberation about careers is anything but fast and frugal, and so my framing the decision in terms of a heuristically-mediated decision may seem misplaced. However, I think instead that we should view this use of a heuristic as what might be called a *slow heuristic*. The presumption in the Raz case (which I accept) is that even given all the time in the world, a normal person would still not be able to decide in advance about whether the career as a philosopher or a lawyer would be best. Nonetheless, we can easily recognize that we are already ignoring many small features in the possible futures when comparing the heuristic state “life-as-a-lawyer” to the heuristic state “life-as-a-philosopher.” And even when there is a distinct “life-as-a-lawyer+” to compare, we nevertheless still use the exact same heuristic state “life-as-a-lawyer” when comparing  $L+$  to “life-as-a-philosopher.” And that is the case even though, when comparing the two lives-as-lawyers just to each other, we can easily make a comparative judgment to show that the  $L+$  is slightly better than  $L$ . (To clarify, I am calling the use of “life-as-a-lawyer” and “life-as-a-philosopher” *slow* heuristics because they have involved a significant amount of thought; the deliberator does not generate and employ them quickly, as is typical for other heuristics such as an *availability* or *recognition* heuristic).

---

<sup>4</sup> This analysis of the Small Improvements Argument is similar to that in Anderson (2015).

This still leaves open the question of whether it is *irrational* to use a slow heuristic in these cases. What I'd like to suggest is a particular explanation of how the small improvement case works. In considering  $L$  and  $L+$ , what are the differences in the choice situations when comparing each to  $P$ ? Why is it reasonable to not change one's heuristic when shifting from  $L$  to  $L+$  in the choice scenario when one is comparing each to  $P$ , but it *is* rational to prefer  $L+$  to  $L$ ?

When one is comparing  $L$  to  $P$ , one is already ignoring many details of both situations; we must suppose that *some* ignoring of details is rational. The difference between  $L$  and  $L+$  is smaller than features of  $L$  that the agent is already (by stipulation *rationally*) disregarding when considering  $L$ . It is plausible to think that what makes it rational to not shift one's heuristic when the opportunity for  $L+$  becomes available is something like the following heuristic principle, which I will call the *Principle of Comparable Disregard (PCD)*:

**PCD:** If, in generating a heuristic judgment, one is already rationally disregarding fine-grained consideration  $C$ , then it is rational to disregard other considerations of equal or lesser weight than  $C$ .

The PCD is a heuristic both in the familiar sense that it is a rule of thumb, but also in the more technical sense that it recommends substituting a simpler attribute (or set of attributes) for a more complex one for the sake of making deliberation more manageable.

I should say that the PCD is *not* universally valid. That's because enough iterations of it could allow a situation to pass a threshold whereby the new consideration *does* make a moral difference. But its point is not to be a universally valid principle. In *most* cases, the PCD is an appropriate heuristic principle to employ when considering options about which one lacks full knowledge, and one must

make a decision on the basis of such incomplete knowledge. So, I want to emphasize the practical role that the PCD is playing. It is (1) a rule of thumb that people likely employ (even if non-explicitly) in the face of having too many considerations to take into account while making a decision in a context; and it is (2) reasonable to employ in virtually all contexts.

There is a long history within utilitarian theory that relates to this issue, though it has typically not been discussed in the vocabulary of *heuristics*. For the utilitarian, individuals ought to maximize the good, but it is not the case that individuals ought to spend their time thinking about how to maximize the good. Utilitarians have long held that utilitarian theory is a theory of right action rather than a decision procedure (see Bales 1971). R.M. Hare famously argues (1981, Part I) for *two-level utilitarianism*, where it is appropriate for individuals to make ordinary decisions at an intuitive level, only going to a critical level when necessary. As Jeremy Bentham writes, “It is not to be expected that this process [of utilitarian calculus] should be strictly pursued previously to every moral judgment... It may, however, be always kept in view: and as near as the process actually pursued on these occasions approaches to it, so near will such process approach to the character of an exact one” (1789, Chap. IV, Sec. VI; see also Sinnott-Armstrong 2021, §4). On perhaps the best recent version of a two-level view, Fred Feldman (2012) gives an account of a utilitarian decision procedure in cases in which one does not know how to maximize utility, and the view here should be seen as fitting within Feldman’s general framework. (Additionally, Yetter Chappell 2019, 105, recommends a consequentialist use of heuristics, and Armendt 2019 does so in defense of causal decision theory.). I will have more to say about this below, but the point for the moment is that principles such as the PCD, which can be used in conjunction with slow heuristics, can provide a bridge between one’s inability to perform a full-fledged expected utility

calculation under conditions of boundedness while still “keeping in view” the general idea of maximizing expected utility.

### ***4.3. An application of a two-level account***

An analogy with cigarette smoking may be helpful in showing the usefulness of the kind of two-level view I have in mind. What is the impact of smoking a *single* cigarette? The causal relationship between cigarette smoking and cancer is small, diffuse, and probabilistic, and still not fully understood. Smoking can cause harms and lower life expectancy through multiple channels, and some people who smoke live long and quite healthy (and cancer-free) lives. The impact of one individual cigarette is tiny (at least for someone who is already a smoker). Is there nothing negative (from the perspective of one’s own long-term self-interest) in smoking any single cigarette?

Kingston and Sinnott-Armstrong (2018) argue that the causal relationship between individual GHG emissions is tiny, diffuse, probabilistic, and in general best seen at levels of explanation higher than the individual level. At the same time, the formation of cancer is arguably emergent (see Plutynski 2020, Ch. 1 for discussion of the complexities in causal attribution of cancer) in a way analogous to how Kingston and Sinnott-Armstrong claim that climate change is emergent; and if so, then on Kingston and Sinnott-Armstrong’s reasoning, smoking individual cigarettes does not cause cancer, and then there would be no reason not to smoke any given individual cigarette. However, this reasoning can be *iterated perpetually*, every time one considers having a cigarette. Of course, smoking is addictive in a way that driving is not (though driving arguably might

positively correlate with driving on later days, as I shall note below), but the point remains that insofar as there is a choice involved in smoking individual cigarettes, there would be no reason to not smoke, *if* it were true that tiny or causally diffuse risks can always be discounted.

On the other hand, while a Kingston/Sinnott-Armstrong style argument might entail that one can reasonably discount individual cigarette smoking risks, and thus show too much, a two-level view does not. It seems reasonable to say, in accord with an expected value approach, that every cigarette increases one's chances of health complications by some small amount; in a poignantly titled research letter ("Time for a Smoke?...") Shaw et al. (2000) argue that every cigarette reduces life expectancy by 11 minutes. Furthermore, for those who smoke, it does not seem that *other* risks are already being ignored in the smoking of an individual cigarette that exceed the health risks of smoking. So, a two-level view, supplemented with PCD, provides no reason to think that the risks of smoking should be ignored in deliberations whether to smoke any particular cigarette. Maybe this is intuitively the right outcome. Or maybe not – perhaps some *other* heuristic considerations can be used to show that it is reasonable to ignore the risk of smoking some individual cigarettes on some special occasions; but whatever *those* heuristic grounds are, given the fact that the risks still aggregate on the negative side of the expected value ledger, the risk of a lifetime of cigarette smoking is not something that it is reasonable to ignore. (I will say more about lifetime decisions in §6.1.)

## **5. How to think about climate risks from individual GHG emissions**



How does the PCD apply in the case of individual actions that emit GHGs? For one thing, in the particular case of driving, we already know that there are risks involved. According to the United Nations, there are 1.3 *million* deaths and 50 million injuries annually from traffic accidents (2020).<sup>5</sup> On the face of it, in choosing to go for a drive, one is already ignoring risk, or at least not letting risk overwhelm other factors. I will say more about this shortly, but one might suppose at the outset that if it is reasonable to ignore the risk of direct road death, then it is also reasonable to ignore the risk of climate-related risk from driving. Perhaps people don't quite *ignore* risks from driving, as people do take some precautions; but still, people choose to drive in the face of risks, and for practical purposes typically don't bother to consider driving risks as playing any role in particular decisions whether or not to drive somewhere. If that's the case, and if, intuitively, climate risks are less significant than driving risks, then PCD would reasonably permit individuals to ignore them as well.

What does the most recent research say about the expected climate impacts of individual GHG-emitting life activities? Broome (2019) argues persuasively that the expected harm from individual GHG emissions is positive, and Broome (2021) uses the data compiled in Carleton et al. (2019) to arrive at average lifelong harms for a person in a developed country. On a low-estimate model in Carleton et al. (2019) of overall expected harm, Broome (2021) argues that individuals are responsible for approximately six *months* of harm to others; on a high model in Carleton et al. (2019), individuals are responsible for six to seven *years* of harm.<sup>6</sup>

---

<sup>5</sup> See Ori 2020 for more on road ethics.

<sup>6</sup> In my own work (Hiller 2011a; 2011b; and 2014), I use an estimate from John Nolt (2011) according to which one individual is responsible for one lifetime's worth of harm; Broome's (2021) estimates now seem plausible to me, though I have some concerns with Broome's analysis, both for reasons I discuss below, and also due to issues of harm to non-humans. Additionally, in Hiller (2011) I begin (349) with the question of how much harm individual GHG-

Although I endorse most of Broome (2019) and Broome (2021), Broome's calculations are not in fact a full employment of expected value theory. Broome discusses the impacts of individual GHG emissions, but Broome's calculation does not determine the *marginal* effects of one's GHG-emitting activities. Just because one's action can be expected to contribute, say, .001% to a harm of  $x$  units, it doesn't follow that one is morally responsible for  $1/100,000x$  units of harm. That's because what matters, according to expected value theory, is not one's personal contribution, but one's expected *marginal* contribution – the expected *difference* that one makes. For instance, if one chooses to not purchase a tank's worth of gas at a gas station, and thus not emit particular carbon atoms into the atmosphere, someone *else* will most likely emit those very carbon atoms. So, the relevant questions are: how much of a difference to overall *net* global emissions is it expected to make when one increases or reduces one's individual emissions by a certain amount?<sup>7</sup> And once that is determined, how much of an expected difference in climate harm does that *difference* make?

This is the kind of issue that often arises in regard to inefficacy arguments in animal ethics. There is a difference between arguments regarding individual GHG emissions and arguments regarding consuming animal products. In the case of eating meat, the animal consumed is already dead, and

---

emitting acts *cause*. Gunnemyr (2019) criticizes the claim that individual GHG-emitting acts *cause* harms, and I agree with much of Gunnemyr's critique. But I wish to emphasize here that for expected value theory it is not *causing* harm, but *expected* harm, that matters. (Also, in Hiller (2011) I misleadingly use the phrase "causes an expected harm" (355); this phrase is infelicitous. One does not *cause* an expected harm with an action. Rather, actions *have* expected harms.)

<sup>7</sup> For a discussion of this point, see Hale 2011. As I argue in Hiller (2011b), the *timing* of emissions matters, so even if an earlier GHG emission is replaced by an equivalent later emission, the two will not be equal in their climate-related effects.

so insofar as the wrongness of eating meat consists in causing harm, the wrongness must somehow be due to the ways in which one's meat purchase has incentivized the future harm to other animals via market mechanisms. On the other hand, the direct causal effect of one's GHG emissions has the potential, at least, to marginally increase global warming. For this reason, Broome believes that it is wrong to emit greenhouse gases – there is a risk that one's very emissions will play a causal role. But it should be noted Broome is explicitly *not* using utilitarian reasoning when discussing individual obligations (see 2021, 290, and 2012, Ch. 4). Broome is concerned with the very particles that one oneself is responsible for emitting, regardless of whether others would have emitted them.

This is relevant, for example, because in the time of measures to restrict movement due to the COVID pandemic, many bemoaned that GHG emissions were not reduced by as much as one would have hoped (see Tollefsen 2021). For example, some airlines continue to fly airplanes even with few or no passengers – so-called “ghost flights” – simply to preserve airport slots,<sup>8</sup> which gives evidence that even if consumers choose to reduce their personal emissions, it will not have an equivalent effect on net emissions reduction (though it should be noted that the percentage of ghost flights is still said to be “minute” relative to overall flights; also see Jiang et al. 2021 for a detailed and not-pessimistic analysis of the relationship between COVID and emissions).

Economists study *demand* and *supply elasticities*; this means, respectively, the changes in consumption, or production, of a product when the price of it goes up or down. In general, fossil

---

<sup>8</sup> See, for example, Chris Stokel-Walker, “Thousands of Planes Are Flying Empty and No One Can Stop Them,” *Wired*, February 2, 2022, <https://www.wired.com/story/airplanes-empty-slots-covid/>.

fuel supply is rather *inelastic* in the short-medium term; there is a long supply chain between extraction of coal or oil and consumption, and short-term changes in demand won't have a direct effect on short-term production. For instance, Green and Denniss (2018, §3.3) discuss the phenomenon of infrastructure "lock-in". Producers make large up-front investments in production capacity, like investing in coal mines or oil fields. Even if overall consumer demand goes down, sending prices below the average level where the overall long-term investment in the infrastructure is profitable, the producer may still continue to produce and sell the product at the low amount. If the producer is already locked into the up-front investment in the infrastructure, continuing to sell the product at that price may still be a current marginal gain.

Furthermore, how an increased supply translates into increased GHG emissions is complex. It may drive prices down, driving consumption up. However, demand, on the whole, is also fairly inelastic relative to price in the short-medium term:<sup>9</sup> this is because people have regular schedules in which they commute to work, make travel plans, etc., and in general do not change their fossil fuel consumption significantly in direct response to price changes. The relative demand inelasticity of fossil fuels means that a scenario implied by Johnson (2003), where some people start forgoing a limited resource, leading to lower prices, which in turn incentivizes others to increase consumption of the resource – would not be a significant concern for fossil fuels. Over the long term, however, both supply and demand are indeed somewhat sensitive to changes in price (see Güntner, 2014, Krichene, 2002), which means that long-term reductions in consumption can indeed be expected to lead to long-term reductions in supply.

---

<sup>9</sup> See Michael Morris, "Gasoline prices tend to have little effect on demand for car travel," *U.S. Energy Information Administration*, December 17, 2014, <https://www.eia.gov/todayinenergy/detail.php?id=19191>.

If one individual refrains from buying a tank of gasoline and emitting the GHGs in it, those GHG molecules will instead be emitted by someone else; but that next person would then not emit what would otherwise have been their own tank's worth of GHGs, which would then be emitted by someone else, and so on. But the first person's restraint means that, at any given point in time, setting aside larger macroeconomic factors, that amount fewer net GHGs will have been emitted. Once we include macroeconomic factors, perhaps the amount of difference in net emissions will be less than the amount that the individual has refrained from emitting, but all in all, there is little reason to believe that over either the short or long run, the difference in overall emissions is insensitive to individual reductions in consumption.

There are a couple upshots of the considerations from the economists and from Broome in this section. First, there is a great deal of uncertainty, still, about the *marginal* effects of individual actions, even lifetime individual actions, and that, despite excellent efforts from Broome, our best estimates may still not represent the true range of possibility of climate-related effects, and that this is a reason for further investigation on the question. Economists are interested in market changes due to changes in consumer activity, and ethicists are interested in expected values of decisions, and while ethicists can look to economists for clues, because the questions are not the same, the empirical work from economists does not always directly address the questions to which ethicists wish for answers. That being said, it seems that because of the inelasticity of demand in the short-medium term, and increased elasticity of supply in the long term, it is not unreasonable to think that the marginal effects of GHG-emitting activities do not depart *dramatically* from the average effects as calculated by Broome (2021).

Finally, we can return to the question with which I began this section: How does the PCD apply to climate risks from driving? According to the United States National Highway Traffic Safety Administration (see undated “Summary Table”), there were approximately 1.11 traffic fatalities per 100 million vehicle miles traveled in the US in 2019. According to the US Environmental Protection Agency (EPA), “The average passenger vehicle emits about 404 grams of CO<sub>2</sub> per mile” (2021), or in other words, 40,400 metric tonnes of CO<sub>2</sub> per 100 million miles. Using Broome’s (2021) lower harm estimate, this amounts to approximately 17 life-years per 100 million miles, and on the higher harm estimate, it amounts to approximately 220 life-years.<sup>10</sup> This suggests that the risk of causing traffic fatalities is just above the expected harm from a car’s GHG emissions in the low-harm estimate, and quite a bit below it in the high-harm estimate.

I confess that these are still all rather rough estimates, and I hope that in time, better estimates will be developed. Furthermore, given the above considerations regarding the difference between the average harm and the marginal harm from GHG emissions, I am unsure exactly how to modify these values in light of market inelasticities, and I encourage the reader to come to their own conclusions in light of the details here. My own sense is that while there is still significant uncertainty about the level of climate-related risk of individual actions, it would nevertheless not be reasonable to argue on the basis of the considerations above that it is *clearly* the case that climate risks from driving are less significant than direct traffic fatalities, and to claim that these

---

<sup>10</sup> The lower estimate in Broome (2021) is half a year of harm per 1200 tonnes; solving for  $\frac{.5}{1200} = \frac{x}{40,400}$  yields  $x \approx 17$ . The higher estimate is six-seven years for the same amount of emissions; solving for  $\frac{6.5}{1200} = \frac{x}{40,400}$  yields  $x \approx 220$ .

risks can be fully disregarded. On the one hand, the point of using a heuristic is to avoid having to do these kinds of calculations in the first place. On the other hand, the information I have provided here can be employed by those who may wish to generate a *slow* heuristic using a general ‘approximate climate risks per action.’ But this kind of heuristic does not seem to give reason, on its own or using the PCD, to *ignore* climate risks, even given all the uncertainties involved.

I’d like to end this section by clarifying several open issues. First, I have focused on PCD as a heuristic principle; perhaps there are *other* heuristic principles that can be reasonably employed that make it reasonable to ignore climate-related risks of driving. The PCD cannot be the only applicable heuristic principle, since it itself makes reference to fine-grained considerations *already* being properly ignored. To that extent, my argument is limited. But I should note that even if there are such other heuristic principles that permit the ignoring of climate risk, the upshot would be different than that of typical inefficacy arguments; rather than agreeing that climate risks make no difference in ordinary actions, the most one could say would be that (within a two-level framework) it is sometimes reasonable to ignore them. (The same might be said for ignoring risks from individual cigarettes.) Second, I should also note, in accord with Ori (2020), that we should perhaps also not be so quick to ignore traffic-related risks from driving. If so, then the case for not ignoring climate risks becomes even stronger. Third, I haven’t discussed *other* ordinary GHG-emitting activities. But my hope is that my discussion here of the economics of climate risks and of heuristics (and PCD in particular) can be modeled to apply to other cases as well.

## **6. Beyond individual actions**

### **6.1 A note on single-action vs. lifespan decision-making and planning**

There are certainly many actions whose GHG-related expected disvalue is so small that the PCD indeed likely applies. For instance, when one boils water for tea, one emits greenhouse gases, but one also runs the risk of burning oneself on the pot or the water itself. What does the view in this paper say about such actions? Furthermore, is there a meta-level at which we must engage in a determination to see if PCD applies? The whole point of PCD is to avoid having to over-calculate in particular circumstances.

To be clear, I accept that there is still some small expected disvalue for such actions to be placed on the negative side of the ledger. However, it is still the case that for many such actions, it is reasonable to ignore the GHG-related effects. (As Sinnott-Armstrong (2005) points out, even breathing emits CO<sub>2</sub>; a view that held that one must constantly keep climate change in mind in each breath is absurd.)

While much of the literature regarding individual climate ethics has viewed GHG-emitting choices as independent events, as Michael Bratman has long pointed out (cf. 1987), our lives are not a series of disconnected choices but rather are structured by plans. Dale Jamieson (2007) also notes that, regarding climate change, utilitarians ought to be virtue ethicists, so as to instill durable pro-environmental character traits. And Marion Hourdequin (2010) and Trevor Hedberg (2018) argue that, on grounds of integrity, if one cares about the environment, one should try to refrain from GHG-emitting activities even if the effects are limited. This also relates to the analogy with smoking: it can be argued that the proper way to view smoking decisions is not as decisions to smoke individual cigarettes, but to purchase packs or cartons, or to quit this month or not to quit.



And it is these decisions that have more impact than individual cigarette choice risk. Perhaps driving is similar, insofar as one chooses whether to have a car, or where to live relative to one's job (although it should be noted these choices are often constrained in one way or another by financial limitations), and these decisions have larger impacts on overall GHG emissions than decisions about whether or not to go on individual Sunday drives. (One may also decide on general hobbies for one's days off – one can choose to be the kind of person who travels somewhere distant most weekends or who does activities close to home). My point here is just that from the perspective of individual decision making with regard to climate risk, there are a number of relevant levels of analysis, and the level of analysis of the individual Sunday drive, while not inappropriate, is perhaps not the most important one when considering the individual risk of our decision-making. In this way, there is some truth to taking a broader perspective like in Kingston and Sinnott-Armstrong (2018) and Gesang (2021) – but that doesn't show that *single* individual actions make no difference.

Taking some time at various stages of one's life is consistent with Bentham's suggestion (and Hare's two-level view) to not always focus on abiding by expected utility theory, but still keep in mind the overall goals of maximizing utility. The best level of assessment for individuals with regard to climate change is at the level of the individual's more general life-plan. Individuals, especially those in wealthy nations who are more than capable of doing so, should structure their lifestyle so as to reduce or limit activities, even small-scale ones, that emit GHGs, though of course in our era it is impossible to eliminate them entirely.

## **6.2. On the relation between individual and group action**

One might wonder whether it is frivolous to once again discuss individual action in the contexts of climate change – one often hears claims that we must overthrow the system and not dwell on little things. I would like to make several points in response. First, as I myself have noted (2011b, 365), creating political change faces some of the same inefficacy concerns as reducing climate change. Second, it should be emphasized that telling people in a public forum that individual decisions make a difference (or do not make a difference) is *not* an individual action – it is a collective action. Third, one sometimes hears claims that individual changes will not *stop* climate change. As the headline of an article in Time Magazine by climate scientist Michael Mann puts it (2019), “Lifestyle Changes Aren’t Enough to Save the Planet.” But the truth is that it is now impossible to *stop* climate change. And there has never been a question of “saving the planet”. What it is not too late to do is incrementally lessen the negative impacts of climate change. The rhetoric of “stopping climate change” or “saving the planet” is inappropriate, and perhaps leads people to reject the expected value approach and its incrementalism. The problem here is with the rhetoric, and not with expected value theory.

Sometimes, ignoring decimal places in a numerical claim is conversationally appropriate, according to Gricean norms of conversation (Grice 1989). Telling someone that it is 9:01:17 is conversationally *worse* in many contexts than just saying that it is nine o’clock. Likewise, telling someone about to go for a short drive that their GHG emissions make no difference may, in some conversational contexts, be more conversationally cooperative than saying that it makes a tiny difference. But if the considerations in this paper are correct, then a *general* practice of telling people that short drives make no difference, across conversational contexts, is not appropriate.

Perhaps the claim that individual emissions make no difference can be used to *flout* Gricean quality norms – of course individual GHG emissions have *some* expected disvalue! – so the implicature goes that the disvalue is so small that we should focus on other things, like changing laws or overthrowing the system.

Those are admirable socio-political goals. And philosophers have long discussed *collective* responsibilities (see, e.g., the essays in Bazargan-Forward and Tollesfen, 2020) – a topic I have not broached here. But we can *both* be mindful of our emissions – sometimes, if the view in this paper is correct, in the back of our minds – and also mindful of the best ways to engage in political and other collective activities to best limit climate change. And we should act upon those intentions, both in our personal and – insofar as the personal is not *already* political – in our public and political lives.

## **7. Conclusions**

To sum up, I have argued that although it is reasonable to ignore some risks, one ought, at various points in one's life, consider the climate impacts of one's lifestyle, and attempt to formulate life-plans in ways that take into account the expected harms of one's GHG emissions. I have argued that *some* small risks may reasonably be ignored because human psychology is bounded not just by the amount of time we have to make decisions but also by the near-impossibility in many cases of acquiring conclusive evidence about the details of future states of affairs that may ensue if one chooses an option. I have further argued (in accord with a long line of two-level ethical theorists) that these claims are not in violation of the spirit of expected value theory, which accepts that

there are costs to calculating risks in particular cases, and thus does not require individuals to do so constantly. Nevertheless, even the reasonable use of heuristics does not seem to permit one to ignore the climate-related risks of individual actions (although the climate risks of certain of one's actions may still be outweighed by their positive expected effects), and does not seem to permit the general condoning of ignoring individual actions' climate risks.

My aim in discussing heuristic principles is to give voice to something correct that may underlie some individual inefficacy concerns – that sometimes, risks may be too insignificant to merit consideration – without undermining the view (such as in Hiller 2011 and Broome 2019) that there is still some amount of expected disutility of ordinary actions due to climate risks (and that this disutility is likely not miniscule). Why does this matter? One might wonder who bears *responsibility* for making changes to help reduce climate risk. Perhaps ordinary individuals, in deciding upon *some* minor actions, may be reasonable in ignoring climate-related risks, but the fact that there is always going to be some amount on the negative side of the ledger due to climate risk means that we as individuals do indeed bear some moral responsibility for mitigating the climate risks we impose upon others (in addition to there being collective responsibility for climate risks as well). And as I have tried to argue, here and elsewhere, these risks are not so small so as to be morally insignificant. Inefficacy arguments leave no room for this assessment.

*So much* of the nitty-gritty of life involves doing little things that make little (expected) difference. One might walk one more block for exercise; one might slightly lower the pitch of one's voice in conversation with a friend in need of calm; one might add a shake of garam masala to the pot of dal one is cooking; one might add a touch more vibrato on a 5th note in a guitar riff one is playing;

one might take a multivitamin; one might carry a sign in a protest line; one might make a kind facial expression to a fellow passenger on the bus.

The general perspective from which the denial of the expected utility approach arises – that small actions make no difference – may fully nullify the significance of these actions and thus potentially leave one no reason for doing them. Perhaps advocates of inefficacy arguments can, for each of these domains, provide independent reasons for claiming that these small actions are appropriate or not, despite their making no difference (or no noticeable one). But the vast heterogeneity of cases in which small-scale actions occur, and to which we intuitively attach value judgments, suggests that a general solution is probably the best one. And the general solution given by expected value theory – that small-scale actions do indeed make small *expected* differences, which then can add up – for (expected) better or (expected) worse – is the most plausible and theoretically elegant way of accounting for the value of these actions, even if expected value theory is nevertheless not always action-*guiding* because of the boundedness of human capacities.<sup>11</sup>

## References

- Adler, Matthew D. (2007). Why De Minimis?. *U of Penn, Inst for Law & Econ Research Paper*, (07-12):7-26.
- Anderson, Jack (2015). Resolving the small improvement argument: a defense of the axiom of completeness. *Erasmus Journal for Philosophy and Economics* 8 (1):24-41.
- Andreou, Chrisoula (2015). Parity, Comparability, and Choice. *Journal of Philosophy* 112 (1):5-22.
- Armendt, Brad (2019). Causal Decision Theory and Decision Instability. *Journal of Philosophy* 116 (5):263-277.
- Balfour, Dylan (2021). Pascal's Mugger Strikes Again. *Utilitas* 33 (1):118-124.

---

<sup>11</sup> Many thanks to Adriana Placani and Stearns Broadhead for extremely helpful comments on an earlier version of this paper.

- Bales, R. E. (1971). "Act-utilitarianism: account of right-making characteristics or decision-making procedures?", *American Philosophical Quarterly*, 8: 257-65.
- Barrington, Mitchell (ms.). Ignoring the Improbable.
- Bazargan-Forward, Saba & Tollefsen, Deborah (eds.) (2020). *The Routledge Handbook of Collective Responsibility*. Routledge.
- Bentham, Jeremy (1789). *An Introduction to the Principles of Morals and Legislation*.
- Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Cambridge: Cambridge, MA: Harvard University Press.
- Broome, John (2019). "Against Denialism." *The Monist* 102:110-129.
- Broome, John (2021). "How Much Harm Does Each of Us Do?" In *Philosophy and Climate Change*. Edited by Mark Budolfson, Tristram McPherson, and David Plunkett, Oxford University Press.
- Carleton, T., M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, A. Jina, R. Kopp, K. McCusker, I. Nath, J. Rising, A. Rode, H. K. Seo, J. Simcock, A. Viaene, J. Yuan, and A. Zhang (2019). "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." BFI Working Paper, August 12. Chicago, IL: Becker Friedman Institute for Economics at the University of Chicago.
- Chang, Ruth (2002). The possibility of parity. *Ethics* 112 (4):659-688.
- Chappell, Richard Yetter (2019). "Fittingness Objections to Consequentialism." In *Consequentialism: New Directions, New Problems?*, edited by Christian Seidel. Oxford University Press.
- Cripps, Elizabeth (2013). *Climate Change and the Moral Agent: Individual Duties in an Interdependent World*. Oxford University Press.
- De Sousa, R. B. (1974). The good and the true. *Mind*, 83(332): 534-551.
- Erickson, P., Lazarus, M. (2018) Would constraining US fossil fuel production affect global CO<sub>2</sub> emissions? A case study of US leasing policy. *Climatic Change* 150: 29-42.
- Feldman, Fred (2012). True and Useful: On the Structure of a Two Level Normative Theory. *Utilitas* 24 (2):151-171.
- Fischer, Bob (2019). *The Ethics of Eating Animals: Usually Bad, Sometimes Wrong, Often Permissible*. New York, NY, USA: Routledge.
- Gesang, Bernward (2017). Climate Change—Do I Make a Difference? *Environmental Ethics* 39 (1):3-19.
- Gesang, Bernward (2021). Utilitarianism and Heuristics. *Journal of Value Inquiry* 55 (4):705-723.
- Gigerenzer, Gerd (2000). *Adaptive thinking: Rationality in the real world*. Oxford University Press.
- Gigerenzer, Gerd (2021). "What Is Bounded Rationality". In Viale, Riccardo ed. *Routledge Handbook of Bounded Rationality*. Routledge.
- Green, F., Denniss, R. (2018). Cutting with both arms of the scissors: the economic and political case for restrictive supply-side climate policies. *Climatic Change* 150, 73-87.
- Grice, H. Paul (1989). Logic and Conversation. In *Studies in the Way of Words*. Harvard University Press. pp. 22-40.
- Gunnemyr, Mattias (2019). Causing Global Warming. *Ethical Theory and Moral Practice* 22 (2):399-424.
- Güntner, J. H. (2014). How do oil producers respond to oil demand shocks? *Energy Economics*, 44: 1-13.
- Hale, Benjamin (2011). Nonrenewable Resources and the Inevitability of Outcomes. *The Monist* 94 (3):369-390.

- Hare, R.M. (1981). *Moral Thinking*. Oxford: Clarendon Press.
- Hedberg, Trevor (2018). Climate Change, Moral Integrity, and Obligations to Reduce Individual Greenhouse Gas Emissions. *Ethics, Policy and Environment* 21 (1):64-80.
- Hiller, Avram (2011a). Morally Significant Effects of Ordinary Individual Actions. *Ethics, Policy and Environment* 14 (1):19-21.
- Hiller, Avram (2011b). Climate Change and Individual Responsibility. *The Monist* 94 (3):349-368.
- Hiller, Avram (2014). A 'Famine, Affluence, and Morality' for Climate Change? *Public Affairs Quarterly* 28(1):19-39.
- Hourdequin, Marion (2010). Climate, Collective Action and Individual Ethical Obligations. *Environmental Values* 19 (4):443-464.
- Jiang, Peng, Fan, Y. V., & Klemeš, J. J. (2021). Impacts of COVID-19 on energy demand and consumption: Challenges, lessons and emerging opportunities. *Applied energy*, 285, 116441.
- Kahneman, Daniel, and Frederick, Shane (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In Gilovich, T., Griffin, D. and Kahneman, D. eds., *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press.
- Kahneman, Daniel (2003). Maps of bounded rationality: Psychology for behavioral economics. *American economic review*, 93(5):1449-1475.
- Kahneman, Daniel (2011). *Thinking Fast and Slow*. Macmillan.
- Kingston, Ewan & Sinnott-Armstrong, Walter (2018). What's Wrong with Joyguzzling? *Ethical Theory and Moral Practice* 21 (1):169-186.
- Kosonen, Petra (2021). Discounting Small Probabilities Solves the Intrapersonal Addition Paradox. *Ethics*, 132: 204-17.
- Krichene, N. (2002). World crude oil and natural gas: a demand and supply model. *Energy economics*, 24(6):557-576.
- Jamieson, Dale (2007). "When Utilitarians Should Be Virtue Theorists", *Utilitas* 19(2): 160-183.
- Johnson, Baylor (2003). "Ethical Obligations in a Tragedy of the Commons", *Environmental Values* 12(3): 271-287.
- Lazarus, M., van Asselt, H. (2018). Fossil fuel supply and climate policy: exploring the road less taken. *Climatic Change* 150, 1-13.
- Lundgren, Björn, & Stefánsson, H. Orri (2020). Against the de minimis principle. *Risk Analysis*, 40(5): 908-914.
- Mann, Michael E. (2019). Lifestyle Changes Aren't Enough to Save the Planet. Here's What Could. *Time Magazine* Vol 194 (11).
- Monton, Bradley (2019). How to Avoid Maximizing Expected Utility. *Philosophers' Imprint*, 19: 1-25.
- Morgan-Knapp, Christopher & Goodman, Charles (2015). Consequentialism, Climate Harm and Individual Obligations. *Ethical Theory and Moral Practice* 18 (1):177-190.
- Nath, Rekha (2021). Individual Responsibility, Large-Scale Harms, and Radical Uncertainty. *The Journal of Ethics* 25 (3):267-291.
- Nefsky, Julia (2012). "Consequentialism and the Problem of Collective Harm: A Reply to Kagan," *Philosophy and Public Affairs* 39: 364-95.
- Nolt, John (2011). "How Harmful Are the Average American's Greenhouse Gas Emissions?" *Ethics, Policy & Environment* 14: 3-10.
- Nye, Howard (2021). Why Should We Try to be Sustainable? Expected Consequences and the Ethics of Making an Indeterminate Difference. In Chelsea Miya, Oliver Rossier & Geoffrey

- Rockwell (eds.), *Right Research: Modelling Sustainable Research Practices in the Anthropocene*. Open Book Publishers. pp. 3-35.
- Ori, Meshi (2020). Why Not Road Ethics? *Theoria* 86 (3):389-412.
- Parfit, Derek (1984). *Reasons and Persons*. Oxford University Press.
- Peterson, Martin (2002). What is a de minimis risk?. *Risk Management*, 4(2): 47-55.
- Plutynski, Anya (2018). *Explaining Cancer: Finding Order in Disorder*. New York, NY, USA: Oxford University Press.
- Raz, Joseph (1986). *The Morality of Freedom*. Oxford University Press.
- Sandberg, Joakim (2011). "My Emissions Make No Difference": Climate Change and the Argument from Inconsequentialism. *Environmental Ethics* 33 (3):229-48.
- Shaw, Mary, Richard Mitchell, and Danny Dorling (2000). Time for a smoke? One cigarette reduces your life by 11 minutes. *BMJ (Clinical research ed.)*, 320(7226): 53.
- Simon, Herbert A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Sinnott-Armstrong, Walter (2021). "Consequentialism", *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2021/entries/consequentialism/>>; accessed April 22, 2022.
- Sinnott-Armstrong, Walter (2005). It's Not My Fault: Global Warming and Individual Moral Obligations. In Walter Sinnott-Armstrong & Richard Howarth (eds.), *Perspectives on Climate Change*. Elsevier: 221-253.
- Smith, Nicholas J. J. (2014). Is Evaluative Compositionality a Requirement of Rationality? *Mind*, 123: 457-502.
- de Staël, Madame (Anne-Louise Germaine). 1845. *De L'Allemagne*. Paris: Didot Frères.
- Sunstein, Cass (2002). *Risk and Reason: Safety, Law, and the Environment*. Cambridge University Press.
- Tollefson, Jeff (2021). COVID curbed carbon emissions in 2020—but not by much. *Nature*, 589(7842), pp.343-343.
- United Nations. "With 1.3 million annual road deaths, UN wants to halve number by 2030." *UN News*. December 3, 2021. <https://news.un.org/en/story/2021/12/1107152>; accessed April 22, 2022.
- United States Environmental Protection Agency. "Greenhouse Gas Emissions from a Typical Passenger Vehicle." Last modified July 21, 2021. <https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle>; accessed April 22, 2022.
- United States National Highway Traffic Safety Administration (undated) "Summary Table", <https://www-fars.nhtsa.dot.gov/Main/index.aspx>; accessed April 22, 2022.