

What's Wrong with Joyguzzling?

Ewan Kingston and Walter Sinnott-Armstrong

This is a post-peer-review, pre-copyedit version of an article published in Ethical Theory and Moral Practice. The final authenticated version is available online at:

<https://doi.org/10.1007/s10677-017-9859-1>.

I – The Issue

Our thesis is that there is no moral requirement—or at least no-one has shown there is one—to refrain from emitting reasonable amounts of greenhouse gases (GHGs) solely in order to enjoy oneself. To concretize this general claim, we will focus on one simple example and argue that there is no moral requirement to refrain from joyriding in a gas-guzzler, which we call joyguzzling¹.

Our thesis is about current circumstances for most people. We remain open to the possibility that circumstances might change so as to generate a moral requirement to refrain from joyguzzling. If a government passes laws against joyguzzling, if a near-universal social norm against joyguzzling emerges, or if a certain person promises not to joyguzzle, then that law, promise or social norm might create a moral requirement to refrain from joyguzzling. Our claim is only about common current circumstances when joyguzzling does not violate any law or break any promise.

Crucially, our thesis is only about moral requirements. We do not deny that it is morally better, best, ideal, or virtuous to refrain from joyguzzling². We also do not deny that there is a *pro tanto* moral reason to refrain from joyguzzling, assuming that not all reasons are requirements. We certainly do not deny that organizations (including governments) have a duty to respond adequately to the problem of climate change or that individuals in many circumstances have a moral duty to push those organizations to meet their duties³. We also make no claims about whether or not joyguzzling might affect an individual's responsibility to promote such political action.

Sinnott-Armstrong (2005a) argued against any moral requirement not to joyguzzle. There have been numerous replies. A recent review found that among the dozens of articles published on

¹ This example is often used in the literature by both proponents and opponents of the view we defend. See (Baatz 2014; Hiller 2011; Jamieson 2014). The example may become less apt as alternatives—in this case, high-powered electric vehicles running on purely renewable electricity—become more widely available.

² Individual “green” actions or lifestyles may also be aesthetically better than non-green alternatives. See Kingston (in preparation).

³ Climate change used to be called global warming, and it might more accurately be called climate degradation, as Jamieson has suggested to one of us in conversation. The name does not matter to our argument.

the topic, “most” hold that “individuals have a duty to take at least some steps to reduce their GHG emissions” (Fragnière 2016, 809). Some of the replies reprise and revise arguments criticized in Sinnott-Armstrong (2005a), and some present new arguments. We will deal with the revised arguments first, followed by the new arguments.

II – Virtues

Some critics argue that we need to cultivate “green” virtues in order to fully understand our responsibilities regarding climate change. For instance, Dale Jamieson argues that preserving the virtue of humility, creating a virtue of mindfulness, and rehabilitating the virtue of temperance are all important tasks for us in the Anthropocene. Most notably, he argues for a virtue of respect for nature that might help us navigate the Anthropocene in a prudent way that allows us self-knowledge and pays homage to the way in which the natural world shapes our own identities (Jamieson 2007). Hourdequin (2010) adds that green activists, in particular, may express more integrity and less hypocrisy if they choose not to joyguzzle. These and related claims are typical of the general field of environmental virtue ethics (Cafaro and Sandler 2005).

Our response is simple: virtues are not moral requirements. Even if integrity, humility, mindfulness, temperance, and respect for nature are all important virtues, and even if it is better or ideal to have these virtues, that still is not enough to show any moral requirement not to joyguzzle. Joyguzzling can fail to be virtuous or good—indeed, it can even be vicious and bad—without violating any moral requirement.

This possibility follows from the nature of moral requirements. Moral requirements differ from other moral reasons in which kinds of reactions are appropriate. Someone who violates a moral requirement without any adequate justification or excuse thereby does something morally wrong and becomes liable to some negative sanction (including moral condemnation, anger, or guilt). Nonetheless, someone who conforms to a moral requirement does not always deserve a reward (such as praise). People who murder and steal deserve condemnation or punishment, but people do not normally deserve praise simply for not killing or stealing. In contrast, someone who gives to a particular charity (such as CARE) does deserve praise, whereas someone who fails to give to that particular charity on a particular occasion does not thereby become liable to a negative sanction, even if they have no adequate justification or excuse. This is especially true if they give plenty to other charities or on other occasions. That is why the moral reason to give to that particular charity on that particular occasion is different from a moral requirement.⁴

⁴ Requirements include perfect but not imperfect obligations and duties, because only perfect obligations and duties can be violated by individual acts. This terminology is appropriate here, because our issue is about individual acts of joyguzzling. Notice also that this distinction between requirements and other reasons can be recognized by consequentialists (Sinnott-Armstrong 2005b), so this terminology does not beg any question against consequentialism.

The crucial point for our response is, then, that an act that is not virtuous or good—or even an act that is vicious and bad—need not always thereby become liable to some negative sanction, such as guilt, anger, condemnation, or punishment. If not, then it does not violate any moral requirement. Hence, it takes more argument to show that the act brings liability to some negative sanction than merely to show that the act is not virtuous or is vicious. That is why our critics fall short when they claim only that joyguzzling is not virtuous or is vicious. That claim changes the topic and fails to address the question about moral requirements that concerns us here.

Some virtue ethicists might reply that, even though virtues are not requirements, requirements can be derived from virtues using bridge principles. Rosalind Hursthouse, for example, claims that “an act is right iff it is what a virtuous agent would do in the circumstances” (Hursthouse 1991, 225).⁵ A parallel principle about moral requirements would claim that an act violates a moral requirement if and only if no virtuous agent would do it in the circumstances. But why should we accept either principle? Indeed, both principles seem problematic in cases of supererogation. Consider, again, donating to a particular charity, such as CARE. Benevolence and generosity seem to be virtues, so a fully virtuous person must be benevolent and generous. It also seems that any benevolent and generous person would donate to CARE in times of need if she believes there is no better use for what is donated. (In this example, if another charity would more credibly help, then the benevolent, virtuous person would give to that one).

If so, no fully virtuous person would fail to donate to CARE in such circumstances. The above bridge principle then implies a moral requirement for everyone to donate to CARE on all such occasions. This derivation conflicts with the common intuition that it is supererogatory rather than morally required to donate to this particular charity on this particular occasion. As we argued above, someone who fails to donate on a particular occasion does not thereby become liable to a negative sanction, especially if they give lots to other charities on other occasions. In such cases, the required bridge principle seems dubious at best.

This point extends to joyguzzling. Even if all fully virtuous people would refrain from joyguzzling, that shows at most that refraining from joyguzzling is morally ideal, but it does not show that refraining from joyguzzling is morally required. We are not required to be fully virtuous in the case of joyguzzling any more than in the case of donating to a particular charity.

Opponents could admit that we are not always required to be virtuous but still insist that sometimes we are. Even though some virtues (such as benevolence and generosity) do not ground moral requirements, maybe other virtues (such as honesty or non-malevolence) do ground moral requirements. Maybe so, but then we still need a separate argument to show that environmental virtues (such as humility, mindfulness, temperance, and integrity) are the kinds of virtues that do ground moral requirements. That relation might hold, for example, when

⁵ Hursthouse adds various qualifications in notes and in future publications, but those qualifications do not affect our point here.

acting contrary to a virtue causes harm (as with malevolence) or violates a moral rule (as with dishonesty). However, to apply this principle to the case of joyguzzling, opponents would then need to show that joyguzzling causes harm or violates a moral rule. This approach thus depends on the kinds of arguments that we are going to criticize in future sections. If these other arguments fail, then the appeal to special virtues also fails.

III – Group Causation

Some critics (including Schinkel (2011) and Schwenkenbecher (2014)) charge Sinnott-Armstrong with the mistake of ignoring the effects of sets or groups of actions. The set of all human emissions will cause harm even if an individual act of joyguzzling by itself causes no harm. However, Sinnott-Armstrong never makes the mistake of denying this obvious truth. Indeed, he explicitly discusses the effects of groups of actions as well as groups of people (in his criticisms of “the group principle”).⁶

The crucial question is whether every individual act is morally wrong whenever it falls into a group of acts that jointly cause harm. Some cases of such harmful groups of acts might seem to support an affirmative answer. In ancient Rome, 23 senators stabbed Caesar. Suppose that no single stab was sufficient to kill him, it took at least 10 or 20 stabs to kill him, and he was still alive until minutes after the final stab. As he lay dying, Caesar reportedly said, “Et tu, Brute.” Brutus could have replied, “My stabbing was not necessary to kill you, because, without me, the 22 others still would have stabbed you, so you still would have died. So my act did not make any difference to your life.” This reply would have been accurate, and some views of causation will see Brutus’ stabbing as not a cause of Caesar’s death. These include accounts in which acts must be themselves necessary or necessary parts of sufficient sets in order to be causal. If Brutus did not cause Caesar’s death, then so much the better for our account, which holds analogously that single joyguzzles do not cause climate change. We must however deal with accounts that claim that Brutus caused Caesar’s death, apparently because his act was one of a group of acts that jointly caused Caesar’s death. Why, critics might ask, isn’t my individual act of joyguzzling a cause of climate change in the same way (because it is part of a larger group that causes harm)?

Note the significant differences between the cases: Brutus and the other Senators conspired and intended to kill Caesar, each directly transferred energy (and a knife blade) into Caesar’s body, causing a separate wound. Those factors—coordination, intention, energy transfer, and individual wounds—can explain why each Senator’s act is seen as causing harm and as morally wrong. In contrast, when millions of people joyguzzle during a year, they do not coordinate, intend to cause harm, or directly transfer energy to any particular target. We do not say that

⁶ Nor does Sinnott-Armstrong commit any of Derek Parfit’s “mistakes of moral mathematics”, since Sinnott-Armstrong explicitly discusses the relevant mistake (2005a, note 23) and shows how Parfit’s crucial example of Harmless Torturers differs from the case of joyguzzling. Parfit himself registers doubts about his original analysis of the “mistake” in the unpublished draft “What we together do”. See also Sandberg (2011).

these factors are necessary to make single acts in a harmful group of actions wrong, but we strongly suspect these factors create the impression of causation in the case of Brutus and other similar cases (including Parfit's 'harmless torturer' example, where a thousand torturers each contribute a few millivolts to a harmful shock). If proponents of the argument from group causation believe that they have example to support the principle without the elements of intention, conspiracy, or direct energy transfer, then we would be happy to consider them. But so far we are not convinced that such examples exist.

IV – The Simple Division Approach

Other critics leave the notion of causation behind and instead claim only that a single joyguzzle *contributes* to the overall problem. Several authors use the language of contribution to refer to the way in which the individual act relates to the problem of climate change as a whole. For example,

[A]n individual drive does not itself cause climate change, but it is certainly a contributing factor, and this contribution is significant...one's own fossil fuel emissions contribute to future disasters that will ... harm people. (Rentmeester 2014, 16-17)

Every single emission contributes to the composite problem, which is made of trillions of tiny emissions. Thus there is no threshold for making a literal difference to the problem: every single atom of greenhouse gas emission saved is a literal improvement. (Lane 2012, 59).

The language of "contribution" allows these writers to link the individual act—indeed, "every single atom"—with the global outcome, without making a firm statement in the more problematic language of causation. Still, the language of "contribution" raises questions of its own.

How much does an individual act of joyguzzling contribute? One common way to calculate contributions can be called the "simple division" approach. For example, John Nolt (2011) multiplies the total expected effect of climate change by the fraction that represents an average American's proportion of the total emissions. He concludes that "the average American causes through his/her [lifetime] greenhouse gas emissions the serious suffering and/or deaths of two future people" (Nolt 2011, 9).

Technical issues aside, the biggest problem is that the language of contribution and the related simple division approach both seem to beg the question. If contributing is *not* a causal notion, we need to know why contributing violates a moral requirement. And if contributing *is* a causal

notion, the statements that a single emission causes even part of the composite problem become much less plausible.⁷

First, suppose contribution is a causal notion. No-one (in this debate) disputes that humanity's overall greenhouse gas emissions are causing climate change, and humanity's overall emissions are made up of emissions caused by individual humans. The IPCC and other scientific bodies tell us this much. But we should not assume that one can find the morally relevant effect of one individual act by dividing the effect of the set of acts it belongs to into chunks that correspond to the effects of an individual act.

Nolt himself admits the difficulty of drawing conclusions directly from his approach. While his formulation is meant to find the level of harm caused by the "average American" he admits: "there is no such thing as the average American" (Nolt 2011, 3). But then he baldly asserts that "the amount of harm done by the average American is not very different from the amount of harm done by you or me" (ibid). He does not explain (a) how an artificial construct can actually cause harm or (b) why the harm that you or I do is similar to the harm that such an artificial construct does. [*Author's note: After publication, we were alerted that Aaron Maltais (2013) makes a similar point to this, in a nuanced discussion of the topic – EK 2/1/2018*].

Next, suppose contribution is not a causal notion. Then we need to know why it is morally relevant. Presumably, the moral relevance of an individual action is meant to be derived by dividing the total moral relevance of the group of acts by the number of the individual acts, in the same way that an average is created by dividing a total value of a set by the number of members in that set. To see the problem with this approach, consider these two cases suggested by Derek Parfit (1984, ch. 3). Say many miners are trapped in a system of mineshafts, and three people in a rescue team are enough to save them. You could join three people to create a rescue team of four who together would save 100 people, or you could go elsewhere and save 10 miners. Although dividing the lives saved by the number of rescuers would mean your "contribution" of lives saved would be increased by 15 if you joined the rescue mission, it seems obvious that in this case you should go solo and save the other 10.

What this shows is that what matters morally is not the effects of groups of acts divided by the number of acts, but the difference that marginal additions will make. We should not expect to find the harm of a single joyguzzle by multiplying its fraction of humanity's emissions by the total harm caused by humanity's emissions. Those who want to say that a joyguzzle is wrong because of the amount of harm it causes need to focus on the effect of a particular emission. We will turn to this approach in the next two sections.

V – Partial Causation

⁷ Others present separate critiques of the simple division approach: Sandler (2011) stresses the difference between average and marginal effects, while Lauren Hartzell-Nicholls (2012) notes the problems with counting climate harms that impact future people.

Our opponent here typically turns to an approach that tries to claim that the actual difference made by one joyguzzle is that it causes part of the harm of climate change. The easiest way to see this would be in the form of an argument.

- (1) Increasing the CO₂ concentration by 100 parts per million (ppm) within decades will cause very large amounts of deaths and suffering S from floods, storms, droughts, sea level rise, and disease transmission.
 - (2) A joyguzzle can be expected to increase the concentration of CO₂ by one part per quintillion ($1/10^{18}$).⁸
 - (3) If CO₂ concentration increases of 100ppm within decades causes suffering S , then increases of 1 part per quintillion causes suffering roughly equal to $S/10^{14}$.⁹
 - (4) Therefore a joyguzzle causes suffering roughly equal to $S/10^{14}$ (from 1-3)
 - (5) $S/10^{14}$ is enough suffering to provide a moral requirement not to joyguzzle.
- Therefore
- (6) There is a moral requirement to refrain from joyguzzling.

We know (1) from the climate science collected in the IPCC reports. Partial causation theorists suggest we should accept (2) since human emissions are already overwhelming the carbon cycle's ability to absorb increases in CO₂, so any additional emissions (however small) must increase the concentration in the atmosphere. We could object to (2) by pointing to the immense complexity of the carbon cycle, which involves the annual cycling of hundreds of billions of tonnes of carbon (most of it from natural sinks and sources) involving complicated feedback mechanisms, and seasonal effects.¹⁰ But here we want to focus on the implausibility of premise 3. We have two arguments against it.¹¹

⁸ Humanity raises the concentration of CO₂ in the atmosphere by around 3 parts per million annually. Grant, for the sake of argument that a joyguzzle raises the concentration in proportion to its contribution to humanity's emissions. Since a 14kg joyguzzle represents around a third of one trillionth (0.35×10^{-12}) of humanity's annual emissions, a single joyguzzle would increase the concentration by around one part per quintillion (10^{18}). This calculation is rough of course, but the precise figure does not affect our argument.

⁹ 100 parts per million is equivalent to $1/10^4$ of the total particles and we saw one part per quintillion is $1/10^{18}$ of the total particles. So the relative proportion of S for a single joyguzzle is $1/10^{18} \div 1/10^4 = 1/10^{14}$

¹⁰ See Jamieson (2014) ch. 6.

¹¹ In addition to the two objections in the text, it is perhaps worth noting that the argument from partial causation depends on the idea of causing part of an effect. That notion might make sense in the case of causing part of (say) a storm, since that might refer to causing a few minutes of the storm or a small increment in the wind-speed or amount of wind or rain. Still, it is not clear why such additional minutes, wind, or rain would count as a harm. It also might make sense to refer to part of a harm such as pain, since that might refer to a few seconds of pain or a small increment in the intensity of the pain. However, one of the central harms at stake in climate change is death, and it is hard to make sense of the notion of part of a death. If a defender of the argument from partial causation replies that a joyguzzle could cause part of the total number of deaths by causing one death, then we would be back to the question of whether a single joyguzzle causes the whole death instead of just part. These issues are more challenging

i) Emergence objection

Premise (3) assumes that increases in the level of CO₂ are harmful in an aggregative way. The effect of the whole is made up of effects of the parts, no matter how small the parts are. It might make sense to determine the effects of a part by dividing the effects of wholes in other cases. Just consider a quantity of oil that has a mass of one kilogram and contains, say, 3 times 10²⁵ molecules of oil. Then we can calculate the mass of one molecule of oil simply by dividing one kilogram by 3 times 10²⁵. The calculation works in this case, because each molecule has a mass, and the total mass is the sum of its parts. This kind of property is called “aggregative” (Wimsatt 2007, chap. 12).

Contrast sliminess. The quart of oil is very slimy, but an individual molecule of oil by itself is not slimy at all. It is not that the molecule has a little sliminess, but much less than the whole quart. An individual molecule is not slimy in the least. We cannot feel any individual molecule at all, so it cannot feel like slime. The same point applies to other properties of the oil, including appearing yellowish and causing moving parts to last longer. This kind of property is called “emergent” because the property of the whole emerges out of parts that lack that property. Emergent properties of the whole are not properties of the parts.

In our view, climate change is emergent in this way. Just as individual molecules of oil do not cause parts of sensations of sliminess (or yellowish color), so individual molecules of greenhouse gas do not cause parts of dangerous climate impacts. Instead, as with the sliminess and color of oil, what increases the dangerous impacts of climate change is larger groups of molecules of greenhouse gases.

Moreover, the effects of these groups do not depend simply on the number of molecules, but also on their arrangement, structure, or organization. The mass of a quantity of oil is the product of the number and mass of its individual molecules; so it has the same mass, regardless of whether it is arranged in a can or instead as a gigantic sheet only one molecule thick. However, it feels slimy and appears yellowish only in the can and not in the sheet. Analogously, the effects of a group of molecules of greenhouse gases also depend on their arrangement. When the molecules are arranged properly in the atmosphere, the group as a whole causes climate change (as well as individual storms), but the same molecules would not cause climate change (or individual storms) if they were re-arranged into a thin sheet only one molecule thick far from the earth’s surface. In this case, any photon absorbed and re-emitted by a particular molecule would most likely be released at one of the many angles that would see it miss the earth, rather than back towards it as typically happens when the molecules are arranged thickly nearer the earth. These analogies suggest that the effects of greenhouse gases on climate change are not like the mass of oil but instead more like the sliminess and color of oil. If so, they are also emergent, and then it is just as inaccurate and misleading to say that

than many assume, but we will not emphasize them here, because our other two objections in the text should be enough.

individual molecules of greenhouse gases increase climate harms as it is to say that individual molecules of oil are slimy or yellowish.

This analogy might seem misleading, because joyguzzling emits not only single molecules but also groups of molecules. Opponents who see a particular case of joyguzzling as raising the risk need not claim that individual CO₂ molecules increase the harms of climate change. All they need to claim is that the small groups of molecules emitted by joyguzzling are enough to cause harms of climate change. However, very small groups of molecules of oil are still not enough to cause feelings of sliminess or yellowish color. Millions of molecules of oil are still too small to feel or see. The analogous question, then, is whether the molecules emitted by a single act of joyguzzling are enough to increase the amount of harm to some human. It is far from obvious that they are enough.

To approach this issue, consider what happens to the greenhouse gas molecules after they are emitted by joyguzzling. They do not continue to cling together as a group. Instead, winds carry them all around the world. That makes it very implausible to claim that these particular molecules as a group could cause any particular climate event or harm, especially because such effects are supposed to occur long after the particular act of joyguzzling, and hence long after the emitted molecules are spread far and wide.

Although we believe that global climate change as well as specific climate events and their harms are all emergent phenomena in the way that we explained, some readers might not be convinced. The existence of emergent properties is controversial in metaphysics, as is the claim that emergent properties can cause any effects. But at least our opponents have not shown that global warming is not emergent. As a result, they have not shown joyguzzling is causal, so they cannot cite partial causation to justify their claim that there is a moral requirement to refrain from joyguzzling.

ii) **Timing Objection**

Thus far, we have used the concept of emergence to argue against premise (3) in the partial causation argument. We have argued that tiny increases in the concentration of CO₂ do not have the same kind of effect as major increases. But even if our argument from emergence fails, we have another response to argument (1)-(6), which would amend the *amount* of suffering attributable to a single joyguzzle, meaning that premise (5), which states that this amount is significant, would become false.

The partial causation approach admits that the wrongfulness of a joyguzzle depends on what difference it makes. To determine what difference it makes, we need to compare what happens after the joyguzzle with what would have happened without the joyguzzle. That, of course, depends on what else would have happened. It seems clear that, even if one joyguzzle does not occur, other joyguzzles as well as other emissions will continue to occur. Humans collectively

are adding carbon to the atmosphere at a rate of more than a hundred tonnes per second.¹² Thus, if a single joyguzzle had not occurred and emitted 14 kg of CO₂, another 14 kg of CO₂ would have been emitted and would have begun to disperse only *a fraction of a second later*, and would have dispersed widely before any harm occurred. This suggests that, even if the partial causation theorists are right about their other assumptions, the overall amount of CO₂ in the atmosphere would have reached the same levels only *a fraction of a second later* than it would have with this particular joyguzzle. [Author's note: after publication, we were alerted to the fact that Aaron Maltais first suggests this line of reasoning in *Maltais (2013)*, but does not develop it in depth – EK 2/1/2018].

The argument from partial causation does not acknowledge that the main effects (if any) of a joyguzzle would simply move the timing of a given total concentration threshold.¹³ Even if we grant, for the sake of argument, that a change in timing of a given concentration threshold did correspond to a change in the timing of the relevant effects, it does not matter whether those impacts happen now or a second later.¹⁴ An injury from a burst seawall, or a death from expanded malaria range, are just as grievous if they happen at time *t* or *t*+1 seconds. Overall, the argument from partial causation glosses over the important matter that joyguzzles primarily affect timing, not absolute temperatures.

The upshot is that even if there were some significant effects of slight changes to the timing of concentration increases, the overall harm of a single joyguzzle cannot be calculated in the way proponents, such as Nolt and Hiller (2011) suggest. The simple division process that they use is bound to return figures (such as Nolt's figure of one American's emissions causing the suffering and death of two people) to be much too high.¹⁵ Because the suffering due to climate change is

¹² Figure extrapolated from data at www.trillionthtonne.org

¹³ This argument does not apply to any effects supposedly caused by the level and timing at which CO₂ levels eventually peak. However, (a) CO₂ levels may not peak in the foreseeable future, (b) bad events caused by the level and timing of the CO₂ peak, if it occurs, are likely to be a very small subset of the total bad effects of global warming.

¹⁴ We think this is a lot to grant - the assumption seems to be that if a joyguzzle does not occur, but another driver emits the same amount of greenhouse gas only a few seconds later, then the overall level of greenhouse gases will reach a threshold a few seconds later and cause some particular harm to occur a few seconds later. We note that this link is implausible, partly because the molecules of greenhouse gas emitted by the other driver are dispersed quickly and widely, perhaps even worldwide. As a result, the fact that they were emitted seconds later does not mean that these molecules will have the same effects seconds later. The process is much too widespread, complex, and chaotic to support that claim. In general, there is no reason to assume that the timing of an emission by seconds affects the timing of the harm much later, as the argument from partial causation would need to assume. This objection applies even if the timing of the harm by seconds were somehow morally relevant.

¹⁵ Nolt's figure of an average American causing the suffering and death of two others via her emissions relies on the guess that climate change will cause 4 billion people to "suffer and/or die" over the next thousand years (Nolt 2011, 9).

not sensitive to split-second timing, the contribution of a single joyguzzle would be much lower, and it is not clearly enough to violate a moral requirement. Thus, even if premise (3) survives the argument from emergence, the timing objection would weaken it considerably, restricting the harms of climate change that a single joyguzzle does to only the very small subset that would be timing-sensitive, if any. Thus the suitably weakened premise (5), which suggests the amount of harm is high enough to be significant, no longer would be plausible.

Our opponents might respond by asking why the figures should make any difference. Is not any harm too much? The problem with this approach, as Dale Jamieson (2014, ch. 5) points out, is that it is too restrictive. If truly minute levels of harm were enough to render something morally impermissible, almost none of our activities would be permissible. In the domain of climate change, cooking food and heating water create some greenhouse gases, and by the partial causation approach, some harm. But if a non-zero amount of harm is enough to create a moral requirement against an action, there would be a moral requirement to forgo greenhouse-gas-intensive “luxuries” such as cooked food and warm water completely. It is hard to see that such a life would morally best, let alone morally required. If the levels of harm don’t matter at all, the partial causation approach proves too much.

Overall, the partial causation argument fails to convince on several crucial points. First, this argument assumes that climate change is aggregative, not emergent, but that assumption is at least unjustified and unclear. Even if a joyguzzle increases CO₂ concentrations at all, we have reason to doubt that increases in CO₂ on the level of parts per quintillion are harmful in the same way that increases of 100 parts per million are. But even if proponents could show that the harmfulness of climate change is aggregative, the main effect of a joyguzzle would be to reach a higher concentration of CO₂ a fraction of a second earlier than it otherwise would have been reached. Since most of the harms we care about are not sensitive to split-second timing, a joyguzzle makes no difference to those harms, and the arithmetic of proponents of partial causation fails. Because of these weaknesses in the argument, we believe the partial causation approach fails.

VI – Expected Disvalue

In response to the problems with the simple division and partial causation approaches, various critics turn to probability and focus on the *expected* disvalue of a joyguzzle, which they assume grounds a moral requirement. The actual value of an act depends only on the values of its actual consequences, whereas its expected value depends also on the values of possible consequences that did not actually occur times the probabilities of those consequences. The expected value from an act often can depart dramatically from its actual value, so this approach might seem to get around the difficult causal pathway from a particular joyguzzle to any harm related to climate change. What’s more, it might seem to involve a much closer connection between the science of probabilistic event attribution, which attempts to measure the increase in likelihood of extreme weather events that can be attributed to climate change.

John Broome is one philosopher who takes the expected disvalue approach: “Each increase in the amount of greenhouse gas in the air slightly increases the quantity of rain, but it will be a matter of chance whether the particular quantity of gas you emit this year will be enough to cause a flood on any particular occasion” (Broome 2012, 76). Morgan-Knapp and Goodman (2015) add the important distinction that the harms of climate change can be divided between extreme weather events, such as the occurrence of particular storms, floods, droughts and heatwaves and gradual changes, such as rising sea levels and expansion of disease ranges. They hold, like Broome, that joyguzzles increase the probability of extreme weather events. They also claim that increases at the joyguzzle level contribute fractionally to the gradual changes. This is essentially the argument from partial causation that we already responded to. But as well as looking at the expected disvalue from extreme weather events, one might hold that since thresholds might exist with regard also to the gradual changes (such as when one small section of land erodes into the sea, taking a family’s house with it), expected-disvalue reasoning plays some role here too. In what follows, we include these threshold effects of gradual changes along with the extreme weather events.

Although a single joyguzzle is very likely to cause no harm at all (not even part of a death), according to the expected-disvalue theorists it is supposed to have a small chance of causing some great harm. What should we make of this argument? The first thing to notice is that it is remarkably similar to the argument from partial causation, which we have already rejected. The argument from expected disvalue looks something like this:

(1′) Increasing the CO₂ concentration by 100 parts per million (ppm) within decades will increase the probability of a set of extreme weather events resulting in expected disvalue D .¹⁶

(2′) A joyguzzle can be expected to raise the concentration of CO₂ by roughly one part per quintillion (10^{18}).¹⁷

(3′) If CO₂ concentration increases of 100ppm within decades have an expected disvalue of D , increases of 1 part per quintillion have expected disvalue roughly equal to $D/10^{14}$.

(4′) Therefore a joyguzzle has expected disvalue roughly equal to $D/10^{14}$ (from 1′-3′)

(5′) $D/10^{14}$ is enough suffering to provide a moral requirement to refrain from joyguzzle.

Therefore:

(6′) There is a moral requirement to refrain from joyguzzling.

i. Emergence

Even if we accept premise (2′) here, the expected disvalue approach requires that the probability of dangerous events can themselves be increased (minutely) by the addition of relatively tiny emissions. But why should we assume this? The Emergence Objection applies just as well to premise 3′ as it does to premise 3 in the argument from partial causation.

¹⁶ If one is a determinist at heart, then some extreme weather events will certainly happen with climate change that would not have happened otherwise, but this is also to say that their probability has been raised from 0 to 1.

¹⁷ For the calculations, see notes 8 and 9 above.

Emergence affects probability as it does other properties. While adding oil to an engine reduces the probability of a moving part failing, it is implausible that adding a molecule of oil reduces that probability of failure by $1/10^{25}$. Adding a molecule to a collection of oil molecules does not increase the likelihood that the collection will appear yellowish. Likewise, we should not assume that a joyguzzle raises the risk of extreme weather events, just because human emissions in total raise the risk of extreme weather events. We have no evidence to lead us to think that the relevant risk-raising property is an aggregative property at this scale. In contrast, the way in which the warming effect of CO₂ depends on the arrangement of CO₂ molecules with each other and the earth gives us a reason to think the property is emergent, so the tiny additions need not increase probability of harms at all.

Proponents of the expected-disvalue argument have not shown us that the ability of large CO₂ increases to *raise the probability* of harmful climate events is an aggregative property, and they need to. Proponents might try to make the idea that the increments of parts per quintillion do have a significant net expected impact plausible by the application of chaos theory (e.g. Morgan Knapp and Goodman 2015). While we do not have the space to discuss interpretations of chaos theory in any depth here, we want to point out that chaos theory does not require that all events at the macro-level of systems (such as the weather) *must* be accounted for by some precise alteration of initial values, it merely notes that this kind of sensitivity is *possible*. And the proponents of the argument from expected disvalue need more than the possibility of chaotic outcomes.¹⁸

Another thought that might be behind the popularity of the expected value approach, despite its implausibility, may be a sorities-style paradox. A 25-mile joyguzzle might emit around 14kg of CO₂. If we admit that a 14kg emission in one afternoon does not raise the probability of climate harms, even by a miniscule amount, then it seems clear that 28 kg emission, consisting of another ineffectual 14kg raises the likelihood no more than the original 14kg, and a 42kg emission raises it no more than a 28kg or 14kg emission and so on. But at some point, adding these 14kg increments, we must get to the kinds of gigaton levels of emissions that do increase the likelihood of extreme weather events. But, as we know, gigaton levels of emissions do increase the likelihood of severe climate impacts. Therefore, 14kg of emissions must raise the likelihood by some amount.

However, this is, as Julia Nefsky (2011) points out, to treat the sorities paradox as if it were a *reductio* of the idea that the addition of no tiny increment (hair, pebble, 14kg of emissions, etc.) can have an effect on the end property (hirsuteness, heapness, risk of climate harms). Simply

¹⁸ Another problem with appeals to chaos might seem to be that increases in CO₂ concentration might prevent some harm, due to butterfly effects, which should count *against* the expected disvalue. Morgan-Knapp and Goodman admit that on their view a tiny particular marginal rise in CO₂ concentrations might prevent a harmful extreme weather event. Still, since there are positive relationships between CO₂ concentrations and temperature and also between temperature and large scale weather events, they assume a marginal rise in CO₂ concentrations is more likely to trigger a chaotic weather event than to prevent it (Morgan-Knapp and Goodman 2015).

noting that eventually, some level of aggregation of tiny increments must sum to a significant increase in the relevant property is not enough to render it obvious that the tiny increment must therefore have some additive property (such as risk-raising) itself.

The difficulty we might face about how large collections of joyguzzle-level emissions need to be to raise the risk of climate harms is not a good reason to assume that even the smallest of changes must have some chance of causing an effect. This assumption would be tantamount to flatly asserting that plucking a single hair has a chance of making someone bald or not. This approach would be wrong in both cases. The intuitive implausibility that a single plucked hair affects a person's baldness is what makes this kind of problem a sorities paradox, rather than a reductio of the premise that a plucked hair makes no difference to a person's baldness. The same approach should be taken with regard to joyguzzles.

ii. Timing

We think that the emergence objection applies to the argument from expected value just as well as it does the argument from partial causation. But, once again, if you are not convinced, there is another reason to distrust the expected value argument – the timing objection.

Recall that if a single joyguzzle affects the concentration of CO₂ at all, its main effect is to increase the timing at which a given concentration of CO₂ is reached by a split-second. Just as with the partial causation argument, the idea of this change raising the likelihood of particular threshold events resulting from climate change becomes far less plausible. Thus, the argument must be amended to remove from the expected disvalue calculation all those events with a probability that is unaffected by such split-second timing changes. This means the total disvalue of a joyguzzle will be a fraction of the extreme weather events, which are themselves a fraction of the effects of climate change. At this point the idea that a joyguzzle must have an expected disvalue that clearly outweighs its cost looks very dubious.

So far we have argued against expected disvalue theorists who hold that, while it is very likely that one's emissions will not cause harm, the risk that they will trigger great harm is enough to justify simple-division approaches for assessing the expected disvalue of a joyguzzle. John Broome (2012) adds an interesting twist to this kind of argument.¹⁹ For Broome, the very long residence time of CO₂ in the atmosphere means that at some point over the centuries, a joyguzzle *will* do harm.

Your emission increases the likelihood of a flood, but it might not actually cause any particular flood. So it is true that your particular emissions may do no harm in a single event. But during the centuries they are in the air they will have the chance of causing harm on innumerable occasions. It is extraordinarily unlikely that they will do no harm at all. There is no real uncertainty there. (Broome 2012: 79)

¹⁹ See also Lawford-Smith (2016) for a similar argument.

Broome himself doesn't see joyguzzling as wrong if we use personal carbon offsets to balance out the harm we do. But the use of such offsets is controversial and we need to respond to his claim that a joyguzzle will, in time, cause harm. Broome's confidence in the harmfulness of individual emissions seems misplaced. While the total effect of humanity's GHG emissions will persist for a significant amount of time, perhaps centuries, the actual atoms of CO₂ emitted by a joyguzzle will likely cycle out of the atmosphere and back into plant matter relatively quickly²⁰. More importantly though, even if my emissions do cause a permanent marginal increase in CO₂, since there are literally billions of acts of emitting greenhouse gases every day, and not billions of harmful weather events, it seems very implausible to say that this increase *will* make the difference between some harmful weather event occurring or not, even over centuries.

Thus far we have shown how it is hard to see how any argument based on causation could establish a moral requirement to refrain from joyguzzling. We have reviewed approaches based on contributions, partial causation, and expected disvalue, and found them all wanting. Since no argument of this kind is sufficient to show that there is a moral requirement, in order to reach the conclusion that there is no such moral requirement, we will need to look at arguments of other kinds, which we assess below.

VII – Fair Shares

The previous replies were to revisions of arguments that were already criticized by Sinnott-Armstrong (2005a). Now we turn to some new kinds of arguments that were not considered in previous defenses of our view.

One attempt to justify a moral requirement not to joyguzzle suggests that to joyguzzle would be to overshoot one's fair share of emissions allowances. Christian Baatz presents a nuanced view of fair shares and states "driving an SUV for fun may not be wrong at all. As long as I stick to my fair share [of rights to emit GHGs], I can drive as much as I please" (Baatz 2014,12). While Baatz here seems to be permissive about joyguzzling, his theory puts the would-be joyguzzler

²⁰ Assessing the "lifetime" of a joyguzzle's emissions is complicated. While there is significant scientific work on the "residence time" of anthropogenic additions to atmospheric CO₂, this focuses on "pulses" of CO₂ on the planetary scale of one to five trillion tonnes, which is at the level of humanity's *total* perturbation to the system over decades (Archer et al. 2009). When we are talking about particular molecules, the residence time is much less clear. As climate scientists David Archer and Victor Brovkin put it, "the CO₂ in the atmosphere [over long periods of time] will not consist of *the exact same* CO₂ molecules emitted from fossil fuel combustion, because of the copious exchange of carbon with the ocean and the land surface" (Archer and Brovkin 2008 [italics mine], 284). Thus the idea that some molecules with my name on them, as it were, remain in the atmosphere potentially causing extreme weather events is a gross oversimplification. It might be that the particular molecules of a single joyguzzle have a counterfactual effect of preventing *some other* CO₂ molecules from being removed from the atmosphere via photosynthesis. But the complexity of even this one natural cycle among the many involved in climate change is one more reason why we should think of climate change harms as emergent

in a difficult position. Fair shares are likely to be very much lower than those used in an average American lifestyle. Currently Americans produce what amounts to 16 tonnes of CO₂ per capita annually, whereas the global average is 5 tonnes, and a sustainable “allowance” will need to be significantly less than that (World Bank n.d.). While one could conceivably use some of one’s precious allowance for joyguzzling, this would probably mean economizing significantly on much more vital activities such as heating, cooling, cooking, bathing, and travelling to see one’s friends and family. Baatz’s approach, in effect, might make joyguzzling personally ill-advised, albeit not morally wrong by itself.

Operationalizing the notion of a fair share of global emissions is difficult. Simon Caney has argued forcefully against the view that each person should be allowed an equal amount of the global sustainable total (Caney 2012). Baatz admits that it is “not possible” to “*specify* to what level individuals must reduce their emissions” (Baatz 2014, 14, italics ours). But what are even approximately fair shares is a tricky question. Compounding the uncertainty is our lack of personal moral responsibility for the socio-economic infrastructure our lives are embedded in, which may require higher emissions for a decent life (Caney 2012). However, in Baatz’s theory, this uncertainty leads not to a total rejection of fair shares but only a limitation on the duties that “fair shares” create. Individuals likely to overshoot a relatively crudely measured fair share only need to take “already available measures” to reduce their emissions “as far as can reasonably be demanded of them” (Baatz 2014, 15), where those actions that can be reasonably demanded are those that are compatible with a minimally decent life. Nonetheless, Baatz argues, on any plausible measurement of fair shares, most of those joyguzzling in the developed world are almost certainly beyond their fair share. This makes it reasonable to demand that these people cease joyguzzling. The upshot is that most cases of joyguzzling at least in the developed world will violate a moral requirement.

Our first response is that the fair shares approach conflates intuitions from two quite different kinds of cases. Compare the following two principles:

- A) When an existing political or social institution encourages cooperation and discourages defection in a collective action problem, individuals have a moral requirement to obey the rules of that existing institution.
- B) When no existing political or social institution adequately encourages cooperation and discourages defection in a collective action problem, individuals have a moral requirement to obey the rules of a hypothetical institution that would be adequate or ideal.

Principle A seems true, while principle B is implausible or at least much more controversial. The force of the fair shares approach comes from, we suggest, the intuitive appeal of principle A. When some good has been formally or culturally apportioned out into fair shares among a society or group, even if in the absence of enforcement, individuals should obey the existing system of division. However, there has been no formally apportioned scheme of fair shares

among individuals regarding joyguzzling or climate change in general. Hence, principle B is the only available basis for a moral requirement to refrain from joyguzzling. Unfortunately, principle B fails in several cases. Imagine a country with an inequitable health-care system, that places the lives of many vulnerable people in jeopardy, increasing inequality. The inequities would be removed if everyone paid 3% more for taxes, and that money funded equitable health care for all. Everyone paying their fair share of an adequate scheme of cooperation would mean individuals devoting a portion of their income to such health-care systems via the government. But in the absence of such a system, there is no basis for a moral requirement on an individual to contribute the same amount of money to the government in the hope that the government would provide such services. Many cases like these refute Principle B (cf. Sinnott-Armstrong 2005a). Since Baatz's version of Fair Shares relies on principle B, this Fair Shares approach cannot establish a moral requirement to refrain from joyguzzling.²¹

VIII – Politics

Even if individuals do not have any moral requirement to forego joyguzzling, they still might gain other moral requirements from their relations to governments. Countries, especially the US, have a moral obligation to help mitigate and adapt to climate change (Caney 2010; Moellendorf 2014; Sinnott-Armstrong 2005a). When countries are failing to meet this moral obligation, individuals might have moral requirements to take steps to “get governments to do their job” (Sinnott-Armstrong 2005a: 312). Baylor Johnson (2003) argues for such a position, drawing on the futility of unilateral action in prisoner's dilemma games. (2013) we might also be required to *create* new political units or organizations to deal with climate change (Cripps 2013; Ostrom 2009).

Here we are talking about joyguzzling instead of political action, so our question is whether a moral requirement to try to get governments to reduce climate change can generate a moral requirement for individuals to refrain from joyguzzling. Several authors claim that certain individual emissions-cutting actions can be appropriate ways to fulfill political duties. Cripps (2013) points out that personal emissions cuts, such as conspicuously and transparently refusing

²¹ In unpublished work, Christian Baatz and Lieske Voget-Kleschin hold that duties to keep to one's fair share are “prior” to any political duties to promote just institutions for dealing with climate change, and thus have moral force on us even in the absence of political institutions regarding climate change. This is because, they hold, (i) we must have a sense of what fair shares are to determine a just institution would be, and (ii) if everyone else were complying with their fair share duty *without* any political institutions, one would still have a duty to stick to one's fair share. Both claims are controversial. Regarding (i), for a way of managing emissions without determining individual fair shares see Caney (2012). Regarding (ii), if our argument above succeeds, if everyone else was refraining from emitting, the duty to stick to one's fair share could be cast as the duty to do what society expects one to do, not to adhere to any “pre-political duty” at all. More importantly, showing that we need a *hypothetical* notion of one's fair share would show nothing about whether that duty applies in the case of widespread non-compliance. [Author's note: Aaron Maltais (2013) also considers, and rejects the fair-shares view of joyguzzling, but for somewhat different reason than we do: the unlikelihood that an effective allocation regime for emissions will be developed in time - EK 2/4/2018].

to fly, can be publicly visible. Lane (2012) adds that individual lifestyle decisions can affect government actions by providing relevant data, such as how many people are spontaneously giving up joyguzzling, that may help a politician support a climate policy. According to Schwenkenbecher (2014), when life-style changes are performed as part of a collective, such as an organized boycott, they might begin to create new collectives that are needed to solve the problem of climate change. Hourdequin (2010) argues that positive personal actions might affect our *own* political awareness and understanding.

These points are all important. Such positive effects can make such lifestyle changes good, ideal, or virtuous. However, such desirable effects are not enough to show that we have any moral requirement to avoid joyguzzling or even to join boycotts or install solar panels. We are not morally required to do everything that is good, since then almost everything we did would be morally wrong for overlooking some better alternative. In order to show a moral requirement to refrain from joyguzzling, we would need to show why political progress cannot be achieved if we ever joyguzzle. That is not shown by any of the arguments so far.

Another group of arguments points to negative effects of joyguzzling, such as hypocrisy or loss of integrity, that might seem to interfere with fulfilling our political duties. Hourdequin writes,

It certainly seems that an individual who worked for emissions limiting policies while steadily and frivolously increasing her own emissions would be working at cross purposes. The kind of unity that integrity recommends requires that an individual work to harmonize her commitments at various levels and achieve a life in which her commitments are embodied not only in a single sphere, but in the various spheres she inhabits. (Hourdequin 2010, 449)

Thus, Hourdequin proposes a moral requirement not to joyguzzle on the grounds that joyguzzling would undermine the integrity of anyone committed to political action to solve climate change. Instead of integrity, Cripps focuses on its opposite, hypocrisy, and notes that avoiding hypocrisy might be especially important for public figures – “those more in the public eye, who are likely to have their actions overhauled by an often unsympathetic press, can least afford to ignore the consideration of perceived hypocrisy” (Cripps 2013, 153). Then, if we have a moral requirement to get governments to do their jobs, this might require public actions that generate an obligation to avoid hypocrisy, which might include refraining from joyguzzling.

However, considerations like these regarding integrity and hypocrisy cannot yield a moral requirement to refrain from joyguzzling. Lifestyle choices can help one fulfill political duties, but refraining from joyguzzling on a single occasion or as one particular lifestyle does not seem necessary or sufficient for meeting fulfilling any political duty. Cripps admits this. She doesn’t go into details in the case of climate change, but she does argue that, while some joyguzzling might need to be foregone to meet one’s political duties, some individual emissions reductions may be unnecessary or even counter-productive to political goals (Cripps 2013, 154).

There seem to be at least three reasons why even climate change activists are not required to give up all joyguzzling. The first is that an individual who fulfills political obligations via lifestyle changes is implicitly or explicitly communicating how much of a lifestyle change affluent people need to make in order to solve climate change. Opinions on this differ wildly. Techno-pessimists think that motorized transport will have to be drastically scaled back to solve climate change, so Sunday driving will never be green. However, techno-optimists believe that advances in electric vehicles together with full decarbonization of the electricity sector will allow us to enjoy Sunday afternoon drives as much as we do today. Some techno-optimists may wish to continue joyguzzling now in order to communicate that this activity need not be sacrificed in the ideal solution, and thereby to allay the fears of a large group of the public.

A second reason is that climate change has become a partisan issue in many countries. An effective political activist may need to avoid distancing herself from powerful groups that oppose policies to fight climate change. If most of her political opponents engage in joyguzzling, then an environmental activist may wish to joyguzzle along with them in order to make it easier to relate to and understand this group's concerns.

Finally, because such lifestyle changes are conspicuous and often unusual, using them as political tools may dilute or confuse the message that our primary duty is to help bring about changes in institutions and policies to solve climate change. Focusing on personal lifestyle changes, even incrementally, risks making the solution seem more about lifestyle than policies. Some individuals might even think that they have done enough by giving up joyguzzling, so this lifestyle change might make them less inclined to do the harder work of politics.

All told, joyguzzling can sometimes conflict with our political duties, but there are also situations in which joyguzzling might not conflict with—and might even help us meet—those political duties. It is probably better in many circumstances for climate change activists to refrain from joyguzzling. It might even be wrong for some actors to joyguzzle in some political communities at some times. Still, none of this shows that we have any general moral requirement to refrain from joyguzzling.²²

This conclusion exemplifies a more general point that it is often risky to move from the political to the personal. A committed climate change activist might feel morally required to speak publicly and lobby government officials for climate change policies. It is not hard to imagine that this activist lives in society where it will be more difficult to fulfill these political goals if the activist gets a facial tattoo. This (perhaps deplorable) situation still does not seem to make it morally wrong for the activist to get a facial tattoo. Personal freedom needs to be protected

²²Some (Hiller 2011a; Jamieson 2014) point to the apparent paradox in our claim that political action might be required but refraining from joyguzzling is not, even though political action arguably has the same emergent-property structure that climate change does. However, Elizabeth Cripps (2013) points out that individual political actions can have a snowball effect, and single actions have even a potential world-changing effect that refraining from joyguzzling does not have. This changes the expected value judgements considerably.

against the encroachment of political obligations in order to allow personal expression and experimentation. Joyguzzling can be seen in this context. If the activist wants to joyguzzle occasionally in spare time on a sunny Sunday afternoon, then the political obligation to work for changes in the government should not be expanded so far as to prevent such personal activities and expressions. We all need to be able to keep our public and private lives separate in ways like this.²³

IX – Conclusion

We have surveyed the main objections to our thesis that there is no moral requirement to refrain from joyguzzling on a particular occasion. This case exemplifies our more general claim that there is no moral requirement to refrain from emitting reasonable amounts of GHGs solely in order to enjoy oneself. Of course, many more objections could be raised. We cannot respond to them all in a reasonable space, but we can continue to respond to more as they arise.

We stress that it may well be optimal or virtuous to refrain from joyguzzling. However, the path to showing a moral *requirement* to refrain from joyguzzling does seem to contain very difficult hurdles. Approaches that try to show an adequate connection between single acts of emitting and the bad effects of climate change must deal with the fiendish complexity of the causal pathways connecting emissions with extreme weather events and gradual harms. Approaches that stress new green virtues will find it hard to justify genuine moral requirements to refrain from emitting rather than *pro tanto* moral reasons to do so. Approaches that focus on the political solutions needed have to show why there is a necessary connection between our political goals, and individually mimicking the behavior that if normalized, would meet the goals. Such approaches also need to guard against the encroachment on personal freedom that a full integration of lifestyle and politics would require. These hurdles are not obviously impossible to jump, but current work has failed to clear them.

Acknowledgments:

We wish to thank Douglas Maclean, John Broome, Christian Baatz, Avram Hiller, Mark Budolfson, Sara Bernstein, Shelly Kagan, Dale Jamieson, and two anonymous reviewers.

REFERENCES

- Archer D, Brovkin V (2008) The millennial atmospheric lifetime of anthropogenic CO₂. *Climatic Change* 90:283-297
- Archer D et al. (2009) Atmospheric lifetime of fossil fuel carbon dioxide. *Annual Rev of Earth and Planetary Sci* 37:117
- Baatz C (2014) Climate change and individual duties to reduce GHG emissions. *Ethics, Policy & Environ* 17:1-19
-

- Broome J (2012) *Climate Matters: Ethics in a Warming World*. WW Norton, New York
- Cafaro P, Sandler R (2005) *Environmental virtue ethics*. Rowman & Littlefield Publishers, Lanham
- Caney S (2010) Climate change and the duties of the advantaged. *Crit Rev of Int Soc and Political Philos* 13:203-228
- Caney S (2012) Just emissions. *Philos & Public Aff* 40:255-300
- Carbon Dioxide Information Analysis Center (n.d) Frequently asked global change questions <http://cdiac.ornl.gov/pns/faq.html>. Accessed 14 March 2017
- Cripps E (2013) *Climate change and the moral agent: individual duties in an interdependent world*. Oxford University Press, Oxford
- Fragnière A (2016) Climate change and individual duties. *Wiley Interdiscip Rev: Climate Change* 7:798-814
- Hartzell-Nichols L (2012) How is climate change harmful? *Ethics & the Environ* 17:97-110
- Hiller A (2011) Climate change and individual responsibility. *The Monist* 94:349-368
- Hourdequin M (2010) Climate, collective action and individual ethical obligations. *Environ Values* 19:443-464
- Hursthouse, R (1991). Virtue theory and abortion. *Philosophy & Public Affairs*, 20:223-246.
- Jamieson D (2007) When utilitarians should be virtue theorists. *Utilitas: A J of Utilitarian Stud* 19:160-183
- Jamieson D (2014) *Reason in a dark time: Why the struggle against climate change failed--and what it means for our future*. Oxford University Press, Oxford
- Kingston E (In preparation) *The aesthetic value of green lifestyles*
- Lane MS (2012) *Eco-republic: What the ancients can teach us about ethics, virtue, and sustainable living* Princeton University Press, Princeton
- Lawford-Smith, H (2016) "Difference-making and individuals' climate-related obligations." In Heyward, C. and Roser, D (eds) *Climate justice in a non-ideal world*. Oxford University Press, Oxford
- Maltais, A (2013) Radically non-ideal climate politics and the obligation to at least vote green. *Environmental Values*, 22: 589-608.
- Moellendorf D (2014) *The moral challenge of dangerous climate change: values, poverty, and policy*. Cambridge University Press, Cambridge
- Morgan-Knapp C, Goodman C (2015) Consequentialism, climate harm and individual obligations *Ethical Theory and Moral Practice* 18:177-190
- Nefsky, J (2011) "Consequentialism and the problem of collective harm: A reply to Kagan." *Philosophy & Public Affairs* 39: 364-395
- Nolt J (2011) How harmful are the average American's greenhouse gas emissions? *Ethics, Policy and Environ* 14:3-10
- Ostrom E (2009) *A polycentric approach for coping with climate change*. World Bank Policy Research Working Paper Series, no. 5095. World Bank, Washington DC
- Rentmeester C (2014) Do no harm: a cross-disciplinary, cross-cultural climate ethics. *De Ethica A J of Philos, Theological and App Ethics* 1:5-21
- Sandberg J (2011) My emissions make no difference. *Environ Ethics* 33:229-248

Sinnott-Armstrong W (2005a) It's not my fault. In: Sinnott-Armstrong, W and Howarth, R Perspectives on climate change: science, economics, politics, ethics. Emerald Group Publishing Limited, New Milford, pp 285-307

Sinnott-Armstrong W (2005b) You ought to be ashamed of yourself (when you violate an imperfect moral obligation)", *Philos Issues* 15:193-208

Wimsatt, W (2007) Re-engineering philosophy for limited beings: Piecewise approximations to reality. Harvard University Press, Cambridge

World Bank (n.d.) World bank open data. <http://data.worldbank.org/>. Accessed 14 March 2017

post-peer review, pre-copyedit version