

# Don't Join the Joyride: Individual Responsibility for Large-Scale Problems

Kjetil Skjerve<sup>a</sup>, Trygve Lavik<sup>b</sup>

<sup>a</sup> Department of Philosophy and Religious Studies, Norwegian University of Science and Technology, Trondheim, Norway, [kjetil.m.skjerve@ntnu.no](mailto:kjetil.m.skjerve@ntnu.no)

<sup>b</sup> University of Bergen, Bergen, Norway, [trygve.lavik@uib.no](mailto:trygve.lavik@uib.no)

DOI: <http://dx.doi.org/10.5324/eip.v13i2.2982>



This is an open access article distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

*The paper argues that, counter to Walter Sinnott-Armstrong and Ewen Kingston's view, we are morally required to refrain from joyguzzling, i.e., driving a fuel-inefficient car for no other purpose than having a good time. It is undisputed that joyguzzling is an example of a situation where the uncoordinated actions of a large group of individuals lead to an undesirable outcome. Additionally, it is highly unlikely that any one individual's actions will have a significant impact on that outcome. But there are morally relevant differences between cases that share these characteristics. The paper clarifies the debate by introducing and discussing three different types of cases: drop-in-the-ocean cases, overkill cases and emergence cases. We argue that we may have moral obligations in drop-in-the-ocean cases, and that emissions of GHGs are not examples of overkill cases. Then we demonstrate through counterexamples that there are moral obligations in a subgroup of emergence cases we call joyguzzling-like cases. After criticizing the soundness of Kingston and Sinnott-Armstrong's arguments, we critically address their relevance. We argue that Sinnott-Armstrong and Kingston fail to distinguish between two concepts of moral obligation — namely, autonomous and heteronomous moral obligation; that their most important arguments do not have any relevance to heteronomous obligations; and, finally, that heteronomous moral obligations are essential for social change.*

**Keywords:** joyguzzling, individual responsibility, climate ethics, harm principle, autonomous vs. heteronomous moral obligations

## Introduction

Walter Sinnott-Armstrong and Ewan Kingston (2018: 311) argue that there is no valid moral principle that offers any justification for an individual moral requirement to refrain from *joyguzzling*. “Joyguzzling” is their term for riding a fuel-inefficient car for no other purpose than to have a good time. Joyguzzling is an example of a situation where uncoordinated actions by a large group of individuals lead to an undesirable outcome, and where it is highly unlikely that one person's actions will have a significant impact on the outcome. The question is whether, and if so why, a

person can have a duty to do something, when it makes little or no difference to the undesirable result whether he or she does it or not.

In the first part of this paper, we clarify the debate by distinguishing between three different types of cases that are often lumped together in the discussion, which we have chosen to call *drop-in-the-ocean cases*, *overkill cases* and *emergence cases*. We argue that drop-in-the-ocean cases clearly raise moral obligations. We are agnostic about moral responsibility in overkill cases, but we argue that joyguzzling is not an overkill case. The force of Sinnott-Armstrong's and Kingston's main arguments depends on understanding joyguzzling as an emergence case. Still, in the second part of the paper, we demonstrate through counterexamples that we do have moral obligations in a subgroup of emergence cases we call *joyguzzling-like cases* (JG cases, for short). While the first and second parts of the paper concern the soundness of Kingston's and Sinnott-Armstrong's argument, the third part addresses its relevance. We argue that Sinnott-Armstrong and Kingston fail to distinguish sufficiently clearly between two concepts of moral obligation – namely, autonomous and heteronomous moral obligation; that their most important arguments do not have any relevance to heteronomous obligations; and finally, that heteronomous moral obligations are essential for social change.

## Some initial premises and clarifications

Like Sinnott-Armstrong and Kingston, we restrict our discussion to whether there are *moral requirements* or *moral obligations* to avoid joyguzzling. Violating a moral requirement deserves blame, while acting according to moral requirements does not necessarily merit praise. Furthermore, denying that A is morally required to refrain from joyguzzling does not imply that it wouldn't be *better* if A did refrain from doing so.<sup>1</sup> Therefore, what we are arguing in the following is not only that it would be *nice* if people refrained from driving cars for fun, but that we may violate a moral obligation by engaging in such driving.

Some works on climate change ethics give detailed principles for distinguishing just from unjust climate emissions (Batz 2014). However, this paper does not aim to address the specifics of which GHG emissions are just and which are not. Instead, our question is whether we can have any obligations at all in cases where our actions cause little or no harm.

We accept the obvious claim that many instances of joyguzzling would collectively cause significant harm. However, this is insufficient in itself to demonstrate a moral requirement to refrain from joyguzzling.<sup>2</sup> We are not entirely convinced that Sinnott-Armstrong and Kingston succeed in showing that *all* the arguments presented in the literature fail. Three arguments come to mind. First, excessive GHG emissions may be in conflict with a virtue, and we may have a moral obligation to be virtuous, regardless of the consequences of our actions.<sup>3</sup> Second, individual actions may serve as good examples, and even if the emissions from a single act are insignificant, our actions may still have a significant impact on other people's behavior.<sup>4</sup> Third, we are not convinced that Sinnott-Armstrong's argument manages to show that the maxim that motivates joyguzzling can be universalized.<sup>5</sup>

We do not find it particularly appealing to walk down any of these paths, since none of them mentions the real reason why we would like people to reduce their carbon footprint. The real reason is not because driving cars for fun turns us into bad

role models, or because it makes us less virtuous, or because the maxim we are following cannot be universalized. Our motive for wanting people to pollute less is, of course, that we would like to minimize the harm caused by global warming. A successful argument for a moral obligation to reduce our carbon footprint must account for this simple fact.

### Three types of cases

We believe there is a morally relevant difference between three kinds of cases where an action contributes to an undesirable outcome: *Drop-in-the-ocean cases*, *overkill cases* and *emergence cases*.<sup>6</sup>

#### a) *Drop-in-the-ocean cases*

*Drop-in-the-ocean cases* are ones where an individual action has a miniscule, but identifiable, morally relevant effect. While the harm or benefit caused by one action may by itself be barely noticeable, the sum of the effects of many actions will be significant. The important characteristic of these cases is that there is a harm or benefit, but not a very significant one. This is to be distinguished from a case where an individual's emissions are too small to cause any harm at all.<sup>7</sup>

Do we have negative moral requirements in drop-in-the-ocean cases? Kingston and Sinnott-Armstrong claim that if it is a moral requirement to refrain from very small emissions that create a "non-zero harm," we have a moral obligation to forgo "cooked food and warm water," which is an implausible position according to Kingston and Sinnott-Armstrong (2018: 177-178).

Kingston and Sinnott-Armstrong seem to assume that a "cutoff line" exists for the moral relevance of harm. However, we cannot see why this would be. For instance, it seems plausible that I am morally obligated to abstain from pinching your finger for no particular reason, even if the harm (let's say around one second of moderate pain) is very small. It is true that any harm, even a small one, would constitute a *reason* for assigning moral duties to individuals. However, a reason for assigning moral duties is not the same as a moral duty. Most people would, for instance, probably agree that animal suffering is a reason for assigning moral duties, but concede that this is compatible with the view that animal suffering is acceptable if it is balanced against sufficiently important human needs. Similarly, small GHG emissions should be balanced against their benefits in order to establish a moral requirement. But Sinnott-Armstrong's main point seems to concern the *harmlessness* of joyguzzling, rather than the *joy* of joyguzzling. Compared to cooking food, the benefits of joyguzzling are inessential and easily replaceable by less polluting activities like having sex, going for a walk, or taking a drive with an electric car. Hence, if joyguzzling is a drop-in-the-ocean case, we do have a moral requirement not to joyguzzle, since even a little harm is sufficient for a negative moral requirement – as long as the action doesn't have any significant benefit.

#### b) *Overkill cases*

In *overkill cases*, the consequences are inevitable, regardless of an agent's actions. The moral question in this type of situation is whether an agent still has a duty to act or

abstain from acting in a way that, under normal circumstances, could have caused the same type of consequence.

Let's say that Chuck, who is hyperallergic to nuts, experiences a strong allergic reaction on an airplane. While a doctor struggles to save Chuck's life, co-passenger Bob continues to enjoy his peanuts because he is convinced (and has every reason to believe) that the damage is already done. Most people will probably feel that Bob is doing something morally reprehensible in this situation, even though he is not harming anyone.

This example shows that it is at least not obvious that we do not have moral obligations in overkill cases. However, we will leave this question aside, since we do not believe that the joyguzzling case is an overkill case.

Counting joyguzzling as an example of overkill would presuppose one of the following assumptions. Either we could assume that we have already passed the tipping point and irreversible feedback mechanisms have already started, in which case any human action will be inconsequential. Or, we could assume that it is politically impossible to make the necessary reductions in GHG emissions and that, for this reason, the tipping point will occur anyway – such as when a wife finally realizes that her husband will never change, so she can safely skip couples therapy without risking any further damage to their already doomed marriage.

Even if we *may* have reached a tipping point, there is not enough scientific consensus on whether this is the case to conclude with a sufficiently high degree of certainty that all hope is lost. The same applies to the claim about politics. We do not know for sure whether it is too late to change climate politics. We cannot rule out the possibility that the majority of citizens in rich countries (for instance because of some disastrous event) will vote for green politicians in the near future in order to make a real change.

Our conclusion is that actions that emit GHGs are not instances of overkill.

### *c) Emergence cases: Structure emergence and threshold emergence*

In drop-in-the-ocean cases, a very harmful effect can be described as the sum of the barely noticeable harmful effects of many actions. By contrast, in *emergence cases* a single action does not by itself cause any harm or benefit at all, but the action belongs to a set of possible actions that together would cause benefit or harm. According to Sinnott-Armstrong's definition, which is close to the standard terminology, a property is emergent if "...the property of the whole emerges out of parts that lack that property" (Kingston and Sinnott-Armstrong 2018: 175).

Since the main object of our inquiry does not primarily concern the properties of objects, but the *effects of actions*, we find it easier to talk about emergence in a slightly different way. According to our more specific application of the concept of emergence, the effects of a set of actions are emergent if the actions that belong to the set do not have those effects individually.

The type of case most closely associated with the word "emergence" is what we choose to call *structural emergence cases*. If A ignites a petrol stove in the tent where C is sleeping, while B closes and seals the vents, this may lead to the effect that C dies from carbon monoxide poisoning. But C would not be "half dead" if either A or B refrained from their action, so the effect that C dies is not the sum of the effects of A's and B's respective actions. In these instances, it is not the "sum" of the effects of the actions that matter, but rather how the actions relate to each other. The beneficial or

harmful effect emerges from the way the actions are structured in relation to one another.

However, the harm to the climate that is caused by human activity does not depend on the structural relation between individual acts. Therefore, we do not believe that climate change is an example of structural emergence. It is the level of emissions that counts, not the relation between the emissions. But there may still be a kind of emergence at play. We believe *threshold emergence* is a more relevant phenomenon in relation to the large-scale effects of individual actions. In *threshold emergence cases* (“threshold cases” for short), the harmful or beneficial effects occur only when the number of similar actions reaches a certain level. In contrast to structure emergence cases, the effects do not rely on any complex relation between the actions in these cases. Let’s say that a mother and father choose not to vaccinate their child against measles. In a community where almost everyone else is vaccinated, the choice does not cause any harm at all. But if the unvaccinated portion of the population exceeds a certain threshold, the community will no longer possess herd immunity, and this may cause serious harm to people who cannot be vaccinated. If the proportion of vaccinated children is already under the critical threshold, the parents’ choice will probably not cause any additional harm (except maybe to their own child). Likewise, if the proportion of vaccinated children is clearly over the critical threshold, the parents’ choice will not cause any additional harm. In the vaccination case, we have no certainty that we are already past the critical threshold at the time of the parents’ decision and that the harm is inevitable. Hence, there is an important contrast between threshold cases and overkill cases. In threshold cases the agent neither knows nor has strong reason to believe that a harmful consequence is inevitable, while the opposite is true for overkill cases.

The feature that makes threshold cases intriguing is that while a set of actions may lead to very serious harm, it is extremely unlikely that any particular action that belongs to the set will make any difference to whether any harm will occur or not. In threshold cases there is an overwhelming probability that the effects of an action are either clearly under or clearly over a tipping point. The question is how a person can be said to cause, or even “contribute to,” something harmful, if it does not make any difference at all whether he or she acts or not.

If joyguzzling is an example of a threshold case, as Sinnott-Armstrong and Kingston seem to believe (Kingston and Sinnott-Armstrong 2018: 176, Sinnott-Armstrong 2005: 299), it is definitely more complex than the previously mentioned vaccination case. There are two reasons for this.

First, there may or may not be a critical threshold where CO<sub>2</sub> emissions reach a “tipping point” that leads to irreversible and radical changes to the environment on a global level.<sup>8</sup> But this is not the only harm threshold. Different thresholds may lead to droughts in specific areas and increased rainfall in others. Some thresholds for increases in sea level will make certain Pacific Islands uninhabitable, while other increases will flood Miami or Copenhagen.

Second, the joyguzzling example may also be a less pure threshold case in the following way. The fact that actions have to reach some threshold to be harmful does not imply that any action beyond that threshold is harmless. For instance, a pollution threshold for the sustainability of coral reefs does not mean that any pollution exceeding this threshold is harmless for the ocean or other ecosystems.

Despite these noted complications, we will simplify matters somewhat for the sake of the argument and assume that joyguzzling is an example of a threshold case. We will furthermore presuppose that a joyguzzling case has the following ethically relevant features:

1. The case is a threshold case.
2. The harm that is caused by the set of similar actions is of grave importance.
3. The agent does not intend for the harmful effect to occur.
4. The agent has knowledge about the likely connection between his or her type of behavior and the harmful effect.
5. The activity is not illegal.
6. The activity is not coordinated.
7. There is a significant possibility that if a sufficient number of people refrained from the activity, the problem would be solved or considerably reduced.
8. It is well known that it would be a good thing if everybody refrained from the activity, and people are frequently encouraged to do so.

Let's call cases with the abovementioned features "joyguzzling cases" – "JG cases" for short – and examine whether we have moral obligations in these cases.

## Moral responsibility in JG cases

### *The harm principle*

Given our assumption that JG cases are a subclass of threshold cases, there is an extremely small chance that any single JG act will cause harm. Sinnott-Armstrong (2005) puts forward a principle he calls "the harm principle": "We have a moral obligation not to perform an act that causes harm to others" (2005: 297). Since a duty (perhaps in contrast to blame or praise) is only relevant to the time of an action and not the time when the consequences of the action have manifested themselves, we take the principle to mean that A has a moral obligation not to perform an act that A *has reason to believe* will cause harm to others.

In Sinnott-Armstrong's wording, the principle states that causing harm is *sufficient* for having a duty not to perform an action. Taken literally, this principle is obviously far too strong. In the drop-in-the-ocean section, we argued that even if harm is always a reason for assigning moral duties, the harm must be balanced with benefits in order to constitute a moral requirement. Hence, a more charitable interpretation of Sinnott-Armstrong would be that harm is always *a reason* for a duty to refrain from action – a reason that can be balanced against other considerations.

Whether or not we violate a moral requirement in JG cases will, of course, depend on what we assume the reasons for a moral obligation actually are. If the harm principle, counter to a literal interpretation, states a *necessary* condition for a moral requirement, we do not have any moral obligations in JG cases. However, if the harm principle is merely taken to state one of several possible *sufficient* conditions for a moral requirement, there may be reasons other than the harmfulness of an act for obligations in JG cases. What makes Sinnott-Armstrong's position compelling seems to be the intuition that harmfulness is the only possible reason for a moral requirement to abstain from something. In the following, then, we will take the harm principle to state a *necessary* condition for negative moral requirements.

## *What does it mean to cause harm?*

There are some pretty obvious cases where we seem to have a duty to refrain from an action, even when the action is neither a necessary nor a sufficient cause of harm. Sinnott-Armstrong (2005: 297) claims that if five people push a car with a passenger locked inside off a cliff, and I join them in pushing the car, my act is neither “necessary nor sufficient to make the car go off the cliff.” This case could probably best be described as an overkill case: a case where the damage is already beyond the control of the agent, but where a person’s action is of a kind that could have contributed to harm under other circumstances. Sinnott-Armstrong (2005: 297) believes as we do that the agent would have a duty to abstain from pushing the car in this case. But he sees the case as an example of causing harm, and his reason is that the action was unusual and the consequence intended.

While we believe this example shows that there are reasons other than causing harm for establishing a moral requirement not to act, Sinnott-Armstrong argues instead for a very inclusive concept of causation. There are two reasons why we do not think this argumentative strategy succeeds. First, Sinnott-Armstrong transforms a relatively straightforward discussion of moral principles into a very complex discussion of the highly opaque concept of causation. With this approach, we will often end up with two difficult questions instead of one. Not only will we have to ask whether the agent caused something (which is a partly empirical and partly conceptual question), we must also ask whether he or she caused the outcome in a *blameworthy* way (which is a moral question).

Second, Sinnott-Armstrong’s alternative criteria for causation – intention and unusualness – are not very robust when it comes to what appear to be irrelevant variations of the case. Let’s say, for instance, that you do not intend to cause harm in the car-pushing case. Rather, you realize that the passenger will die anyway – and even if you are indifferent to his death, you find it amusing to push the car. The lack of intention to kill seems irrelevant to whether you cause harm in this case. Similarly, if joyguzzling does not cause harm, it is not because the action is usual and the harm unintended. Let’s say that Arthur intends to harm the environment because he hates humanity. He installs a large petrol engine in his backyard, and runs it for an hour every Sunday. Even if this action is both intentional and unusual, it would be absurd to say that Arthur’s actions cause climate change while the joyguzzler, emitting the same amount of GHGs, does not.

## *The difference principle*

While it is only implicitly suggested in “It’s Not My Fault” (Sinnott-Armstrong 2005) that the harm principle should be taken to state a necessary condition for moral responsibility, Sinnott-Armstrong and Kingston (2018) seem to commit more clearly to this position in “What’s Wrong.” When they (rightly, in our view) criticize the highly simplistic view that a person is morally responsible for a certain share of the outcome of the set of actions to which his or her action belongs, they suggest an alternative principle: “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” (Kingston and Sinnott-Armstrong 2018: 174). The first part of the sentence corresponds to what we will call “the simple division principle”, and the latter “the difference principle”. We believe that the “simple division principle” – the idea that

we are causally responsible for a certain share of the outcome of the set our action belongs to – is implausible.<sup>9</sup>

The difference principle constitutes a very strong interpretation of the harm principle. According to the difference principle, the only morally relevant aspect of an act is whether the act *makes any difference* to which harmful or beneficial effects occur. If the difference principle is taken to apply universally, neither the cynical airplane passenger nor the car-pusher has done anything morally wrong, since neither of the actions made any difference to the harmful outcomes. As previously mentioned, we agree with Sinnott-Armstrong and Kingston that the division principle is not very convincing. However, we also find the results of the universal application of the difference principle to have some very counterintuitive consequences. Sinnott-Armstrong and Kingston suggest some alternative criteria for moral requirements that are supposed to avoid these consequences.

### *Alternative criteria for moral requirements*

In addition to intention and unusualness, Kingston and Sinnott-Armstrong (2018: 173; Sinnott-Armstrong 2005: 297) mention conspiracy and “direct energy transfer” as some morally relevant aspects of an action. We have already concluded that unusualness and intention are implausible criteria for *causation*. But could these aspects of an action still be relevant to whether we have moral obligations?

What the features above have in common is that some of the most obvious counterexamples to the difference principle have these features, while joyguzzling lacks them. Joyguzzlers do not intend to destroy our planet, the act of joyguzzling is quite common, and joyguzzlers do not consciously participate in a coordinated conspiracy. Further, the causal link between driving a car and the harmful effects of CO<sub>2</sub> emissions is very complex, and not an instance of “direct energy transfer.”

It is not clear whether the conditions are meant as necessary or jointly sufficient criteria for moral responsibility, or maybe as parts of a more complex picture. Regardless, if *some* types of exceptions to the difference principle exist, there is no obvious reason why additional types of exceptions should not also exist. JG cases may very well be one of these types of exception.

One group of counterexamples to the difference principle are cases of complicity, and the aforementioned car-pushing case is of this type. The sixth person who joins the group that pushes the car off a cliff with a passenger inside does not make any difference to the outcome, but he is still undoubtedly both morally and legally complicit in the murder. Is the sixth car-pusher morally responsible because he intended to harm the passenger, and because his act was unusual?

### *Criterion 1: Unusualness*

It is not entirely clear what Sinnott-Armstrong means by the term “unusual.” In some paragraphs he uses the term to refer to occurrences that break with the status quo, while in other paragraphs the term is used “in the sense that most people would not act that way” (2005: 298). Sinnott-Armstrong gives the example of striking a match as an instance of an unusual act. This is understood in contrast to the presence of oxygen, which is not unusual. Here he clearly uses the term “unusual” in the first and not the second sense: striking a match represents a break with the status quo in a way that the presence of oxygen does not. The act of striking the match is therefore

causally responsible for the fire, while the presence of oxygen is not. Still, a lot of people use matches, so striking a match is not unusual in the second sense. While joyguzzling is not unusual in the second sense, since a lot of people in certain parts of the world drive their cars for leisure, it is unusual in the sense that any act of joyguzzling breaks with the status quo. While the conventions at the time and place of an act may not be entirely irrelevant to its moral quality, the claim that “everybody does it” is, in our view, a pretty weak excuse for joyguzzling. You could use the same excuse for supporting slavery in the American South in the early nineteenth century, being a member of the Nazi party in Germany in the thirties or bullying that chubby kid who is bad at sports, as long as bullying kids for these reasons is common.

### *Criterion 2: Intention*

Given the above, we do not agree with Sinnott-Armstrong that unusualness, understood as divergence from what most people do, is an important criterion for moral wrongdoing. Now, what about intention? Are the only exceptions to the difference principle cases where the agent intends to cause harm? What a person intends to do is undoubtedly morally relevant. Many philosophers have also argued that the distinction between intending to do something harmful and foreseeing the harmful side effects of an action is ethically relevant. According to this view, an action executed with the intention of causing harm is morally worse than one in which the agent merely foresees harm being caused, everything else being equal.<sup>10</sup> An example of this kind of reasoning is the argument that intentionally killing innocent civilians in a terrorist attack is worse than unintentionally killing innocent civilians while bombing a hospital where terrorists are hiding. This is a perfectly respectable normative position, but Sinnott-Armstrong seems to imply a much stronger view in his discussion of the complicity case. In contrast to pushing cars from cliffs, you do not intend to cause harm when you ride cars for fun. But as long as you are not a climate change denier, you can foresee that if many people do what you do, it will lead to severe harm. One of Sinnott-Armstrong’s reasons for believing that it is not wrong to joyguzzle is that the joyguzzler does not intend to harm the environment. We believe that this line of reasoning puts far too much emphasis on the difference between intending harm and foreseeing harm as a consequence of actions. Though intentionally causing harm may be morally worse than acting in a way that has foreseeable harmful consequences, both can still be quite bad.

The following complicity case shows that foreseeing the bad consequences of a set of actions that your action belongs to may be sufficient for moral guilt – even when your action is not unusual, there is no direct energy transfer, and your action does not make any difference to the outcome. Walter and Jesse come into your drugstore and want to buy all your packages of the non-prescription nasal spray Sudafed. When you see their car in the parking lot, you have the impression that the two men are going from drugstore to drugstore to buy large quantities of the drug. Your knowledge of chemistry, together with Jesse’s unkempt appearance, lead you to believe that your customers are in the process of obtaining ingredients to produce the illegal narcotic methamphetamine. You could refuse to sell the drug, but you have reason to believe that the two men would then just go to the next drugstore to get what they want. Let’s say you choose to sell the drug, possibly because you do not want to lose the sale, or because you do not want to experience the inconvenience of refusing Walter and Jesse’s request. Selling nasal spray is a pretty common activity for

a pharmacist. You do not intend to contribute to the production of methamphetamine. You do not have any reason to believe that your action will make any difference to the outcome. There is no “direct energy transfer” from your action to the production of the illegal drug. Still, we think it is fairly clear in this example that the pharmacist is responsible for contributing to meth production – and given that meth production constitutes a serious harm, we believe that the druggist has a moral duty to refrain from selling Sudafed in this situation. In some legal systems, such as the Norwegian one, you could probably also make a case that the pharmacist should be charged with criminal complicity.<sup>11</sup> If we assume that there is no law against selling several packages of Sudafed over the counter, the example seems to tick off all the listed criteria for a JG case. Given that the druggist violates a moral obligation in the Sudafed case, we can conclude that we have moral obligations in at least some cases that share all the features of joyguzzling that Kingston and Sinnott-Armstrong point to as ethically relevant.

### *The force of example-driven arguments*

What sort of argumentative force does the Sudafed example have? If you want to make the case that it is morally permissible to joyguzzle, there are at least two possible ways to respond to this case. You could hold your ground and say that the pharmacist had no moral duty to refrain from selling the drug. Or you could come up with some feature that distinguishes the Sudafed example from other JG cases. These two strategies can be used against virtually any example, but none of these responses to the Sudafed case are particularly promising. The only plausible reason to deny that the pharmacist has a duty to refrain from selling the drugs is probably that the consequences of the set of actions are simply not serious enough. But the example can easily be tweaked by stipulating consequences that are uncontroversially serious, which will make it even harder to deny that the druggist has a duty. Let’s say, for instance, that the pharmacist had reason to believe that a customer was driving from drugstore to drugstore in an attempt to obtain enough tranquilizers to commit suicide (and that it was impossible for the pharmacist to intervene in other ways). GHG emissions undoubtedly have extremely serious consequences, so this kind of tweaking would make the example resemble joyguzzling even more closely.

Coming up with a feature that distinguishes joyguzzling from other examples of JG cases also seems to be an unattractive strategy. Adding new criteria – manufactured to separate joyguzzling from cases where we have moral responsibility – does not have any argumentative value, unless there are independent reasons for adapting these criteria.

We have shown that at least some examples of JG cases would indicate that we have moral responsibilities. But just as importantly, we do not think there are examples of JG cases where we clearly do not have any moral duties. If no moral requirement exists to refrain from joyguzzling, we would anticipate that it should be possible to come up with examples that have the same morally relevant features as joyguzzling, but where it is much clearer that we do not have any moral duties. The closest thing to such an example in the two Sinnott-Armstrong articles is the following: If everyone in an airport is talking loudly, some people will miss their flights since they cannot hear the announcements. However, if there is no rule against loud speech, it is not immoral for me to talk loudly since “the harm is going to occur anyway” (2005: 306).

We agree that you do not have a duty to refrain from speaking in this example. But how much does the case resemble the JG cases? The example is obviously a threshold case: the agent does not intend the harmful effect to occur; he or she has knowledge about the likely connection between the type of behavior and the harmful effect; talking at the airport is not illegal; the noisy people are not coordinating their behavior; and if everybody, or a sufficient number, stopped talking loudly, the problem would be solved or considerably reduced. But we find two important differences between JG cases and the airport case. First, people are not encouraged to be quiet in the airport example, while we are constantly encouraged to reduce our CO<sub>2</sub> emissions. Second, the harmful effects of CO<sub>2</sub> emissions are very severe, whereas the harmful effects of not hearing airport announcements are less severe. To illustrate the importance of these points, we can tweak the example somewhat. Let's say the airport had signs every 20 feet that said, "Please keep your voice down to ensure that the announcements can be heard." In this case it would be morally questionable to continue talking loudly just because a large number of other people are doing so. We also believe that the intuitions about the example change if the consequences of the behavior are more severe. Let's say that one of the electric shuttle carts at the airport crashes into a group of people, and that several people are critically injured. Furthermore, imagine that the paramedics are unable to communicate effectively because so many passengers are talking loudly. In this situation the following line of reasoning would be quite preposterous: "The paramedics need people to keep quiet to be able to do their job saving lives. But there are many other people who are not keeping quiet here, so I might as well keep on talking, since it won't make any difference anyway."

Since there are examples of JG cases where it is reasonable to conclude that we have moral obligations, and since no obvious examples of JG cases exist where it is reasonable to conclude that we do not have moral obligations, we can assume that there are moral obligations in all JG cases.

## Two concepts of moral requirements: Autonomous and heteronomous obligations

We have argued that the joyguzzling example belongs to a class of cases where we have a moral duty to refrain from action. This shows that what we have called "the difference principle" does not hold up, even if supplemented with exceptions for unusual or coordinated behavior or for intended consequences. In the following, however, we will argue that even if you are unconvinced by our argument above and believe that the only morally relevant aspect of an action is the (expected) difference the action makes as to whether harmful or beneficial consequences occur, this proves far less than one might initially think. To show this, it is necessary to distinguish between two concepts of moral obligations:

- 1) *Autonomous moral obligation*: An individual is autonomously morally obligated if and only if he or she has a sufficient moral reason for acting in accordance with a norm.
- 2) *Heteronomous moral obligation*: A person is heteronomously obligated by a norm if the community the person belongs to has a sufficient moral reason to sanction the norm by social control. Even if an individual does not have

sufficient moral reasons for acting according to a norm N, society may have moral reasons for wanting people to act according to N, because acting according to N promotes justice. Society may implement norms via social control in the form of praise, condemnation, or critique, and the norms may be internalized as feelings such as anger, guilt, shame, or pride.

While the first concept of moral obligations is presupposed, explicitly or tacitly, in much of the philosophical literature, the second concept is probably closer to a layperson's understanding of what a moral obligation is. If you ask a non-philosopher to identify the current moral norms among people in group G, he or she will probably point to the norms that are socially sanctioned in G, or the norms people in G have internalized – not the norms that individuals in G believe they have rational moral reason for acting in accordance with.

The distinction between heteronomous and autonomous moral obligation should not be conflated with the descriptive-normative distinction. While it is a matter of descriptive fact that a moral rule *is* socially sanctioned, it is a normative question which moral rules *should* be socially sanctioned.

Still, Sinnott-Armstrong and Kingston apparently regard social norms as a subject outside the realm of normative philosophy. They try to take social norms off the table in the following passage: “[I]f a near-universal social norm against joyguzzling emerges (...) [it] might create a moral requirement to refrain from joyguzzling. Our claim is only about common circumstances when joyguzzling does not violate any [social] law” (2018: 170). However, a social norm against joyguzzling does not just suddenly appear. It must be promoted by social sanctions. And it is certainly a normative philosophical question whether social norms against joyguzzling *should* be sanctioned.

Sinnott-Armstrong's and Kingston's strongest case against a moral obligation to abstain from joyguzzling hinges on an autonomous concept of moral obligation. The fact that a single act of joyguzzling causes little or no harm may be significant when the question is whether we have an autonomous moral obligation to refrain from joyguzzling, but it is irrelevant to whether we have heteronomous moral obligations. When the question is whether we should sanction a moral norm against a type of behavior, we are not concerned with a single act; rather, we are concerned with a set of actions. And there is no doubt that joyguzzling belongs to a set of actions that cause severe harm. So, are Sinnott-Armstrong and Kingston denying that we have an autonomous or a heteronomous moral obligation to refrain from joyguzzling?

The answer to this question is not as clear as one might think. Sinnott-Armstrong's explication of the concept of moral requirements fluctuates between the autonomous and the heteronomous perspective: “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)” (2018: 170). The first part of the sentence concerns the person's moral justification for his or her actions, while the second part concerns whether sanctions such as condemnation are appropriate. But, as we have shown, negative sanctions may be appropriate even if the individual does not have any moral reasons to refrain from a specific action. While the above quote is ambiguous, Sinnott-Armstrong is clearly concerned with a heteronomous concept of moral obligation in some other passages. Take, for example, the following sentence: “[B]laming an agent or condemning his act is normally counterproductive when the agent is acting no worse

than other people” (Sinnott-Armstrong 2005: 335). This statement is not about an individual’s moral justification for actions, but a community’s moral justification for socially sanctioning actions. The fact that even Sinnott-Armstrong cannot adhere strictly to an autonomous concept of moral obligations illustrates how limited this concept is.

The duty to vote is another example of how limited an autonomous concept of moral obligations actually is in our context. The central arguments that Sinnott-Armstrong and Kingston use against any moral requirement to refrain from joyguzzling can also be applied to voting. If all those arguments are sound, we also do not have a moral obligation to vote. Hence, when Kingston and Sinnott-Armstrong claim that “individuals might have moral requirements to take steps to ‘get governments to do their job’” (2018: 183), they will run into the same kind of “joyguzzling problems” if they include going to the polls as such “steps.” Voting can only be justified as a *heteronomous moral obligation*.<sup>12</sup> The duty to vote is an explicit social norm in our society, which many of us have internalized. We feel that we are good citizens when we go to the polls. Do we have a similar heteronomous moral obligation to refrain from joyguzzling?

Such a duty must rest on two premises:

- (1) We must have reason to believe that the common good or justice is served if private citizens as a group emit significantly less of GHG gases.
- (2) We must have reason to believe that social control may be an effective tool in the pursuit of changing behavior in this case.

For now, we will assume that the first premise holds. We can justify a social norm against joyguzzling in reference both to the common good and justice. All of us would be better off if no one joyguzzled, and it is unjust that a relatively small group of people harm the environment for everyone else.

Do we have reason to believe that social control can change behavior significantly in our case? Sinnott-Armstrong and Kingston are right that there are no (strong) social norms against joyguzzling or other forms of luxury emissions in most industrialized countries today. If a couple in your neighborhood tell you that they will spend the weekend in Barcelona, you will probably wish them a nice trip, and you would not condemn them for taking a long flight just for a two-day stay. There is a stark contrast here to other examples of unwanted behavior, such as racism. If the same couple express overt racist attitudes, they will most likely be met with condemnation and social exclusion. But what is culturally acceptable can change abruptly and need not be enforced by political actions. The Me Too movement is a good example. What is considered culturally acceptable behavior of men towards women has changed dramatically in a very short period of time, and this change did not come about primarily because of political actions, but because certain acts were called out and socially sanctioned by a handful of people. This indicates that social control can be an effective tool in the pursuit of changing attitudes and behavior.

We have made a case in this section that we are morally required to refrain from activities that lead to excessive GHG emissions in the sense that activities of this kind should be met with negative social sanctions. But how important are heteronomous moral norms in relation to the climate challenge? Is it a deflection to focus on moral norms rather than changing legal norms? One important function of heteronomous moral norms is that they are essential for enforcing legal norms. Consider an example

where corruption is a significant problem in country X. The majority of people recognize the problem and believe that X would be better off without corruption. However, individuals who live there continue to engage in corruption whenever they can, and they justify their individual actions by the same kinds of arguments that Sinnott-Armstrong offers against individual moral obligations regarding the climate problem. Experience tells us that it is almost impossible to fight corruption in a nation if its citizens do not believe there is a moral requirement not to be corrupt.<sup>13</sup> As long as corruption is acceptable on a cultural level, it is almost impossible to fight the problem on a political level. Countries without corruption problems have social norms against doing so. If someone tries to get others to participate in corruption in one of these countries, they will normally be met with moral condemnation. Similar to the corruption case, political attempts to meet the climate problem also depend on social sanctioning of moral norms. Even modest political attempts to restrain private emissions are often met with protests and sabotage. The internalization of a sense of moral duty is an important precondition if we want governments to be able to “do their job.”

## Notes

<sup>1</sup> Sinnott-Armstrong (2005: 296), and Kingston and Sinnott-Armstrong (2018: 170) restrict their discussion to moral requirements, but they argue that virtue ethics cannot solve the problem of whether we have a moral obligation to refrain from joyguzzling, since “virtues are not moral requirements” (Kingston and Sinnott-Armstrong 2018: 170).

<sup>2</sup> See, for instance, Schwenkenbecher (2014: 177), and Sinnott-Armstrong’s and Kingston’s response (2018: 172). The joyguzzling case is not an example of what Derek Parfit calls the second mistake in moral mathematics, that is “Ignoring the Effects of Acts” (Parfit 184: 70-73). Parfit’s examples include coordinated collective actions, which joyguzzling is not.

<sup>3</sup> Kingston and Sinnott-Armstrong claim that “virtues are not moral requirements” (2018: 170).

<sup>4</sup> See, for instance, Schwenkenbecher (2014: 178-179) and the reply in Sinnott-Armstrong and Kingston (2018: 183-184).

<sup>5</sup> For his argument regarding universalization, see Sinnott-Armstrong 2005: 302-303.

<sup>6</sup> Sinnott-Armstrong and Kingston do not make these distinctions explicitly, but implicitly. See Sinnott-Armstrong 2005: 299, 179-180.

<sup>7</sup> See Fragnière (2016: 800), for this other kind of case. Furthermore, we do not think that Derek Parfit’s fourth and fifth mistakes in moral mathematics are relevant for our discussion here. Parfit believes it is a mistake to ignore small effects, if we believe that “enough of [the individuals] will not act in this harmful way (1984: 77, 81). The problem is that if we as individuals refrain from joyguzzling, we cannot assume that enough other car drivers will do the same.

<sup>8</sup> IPCC 2014.

<sup>9</sup> Nolt (2011) and Broom (2012) champion the simple division principle, but we agree with Sinnott-Armstrong’s and Kingston’s criticism of this view (2018: 173ff). We believe that the simple division approach is similar to Derek Parfit’s The Share-of-the-Total view, and we accept Parfit’s criticism of this principle (1984: 67-70).

<sup>10</sup> See, for instance, Anscombe (1958: 12).

<sup>11</sup> See the Norwegian penal code “straffeloven” §§ 15, 22 b and 232.

<sup>12</sup> As far as we know, no one has argued convincingly that casting one vote has anything but a minimal and only theoretical effect on the outcome of an election. We will not discuss Parfit's argument against ignoring small changes further in this paper; however, we notice that Parfit (1984: 74) calculates the expected utility of casting one vote as dramatically higher than similar calculations by Edlin, Gelman and Kaplan (2007).

<sup>13</sup> See, e.g., Fisman and Miguel (2007).

## References

- Anscombe, G. E. M. (1958). Modern moral philosophy. *Philosophy*, 33: 1-19. <https://doi.org/10.1017/S0031819100037943>
- Baatz, C. (2014). Climate change and individual duties to reduce GHG emissions. *Ethics, Policy & Environment* 17: 1–19. <https://doi.org/10.1080/21550085.2014.885406>
- Broome, J. (2012). *Climate matters: ethics in a warming world*, New York: W. W. Norton
- Cohen, G.A. (2008). *Rescuing Justice and equality*, Cambridge, MA: Harvard university press. <https://doi.org/10.4159/9780674029651>
- Fisman, R., Miguel, E. (2007). Corruption, norms, and legal enforcement: Evidence from diplomatic parking tickets. *Journal of Political economy* 115.6: 1020-1048. <https://doi.org/10.1086/527495>
- Edlin, A., Gelman A., Kaplan N. (2007). Voting as a Rational Choice: Why and How People Vote to Improve the Well-Being of Others, *Rationality and Society*, 9: 219–314. <https://doi.org/10.3386/w13562>
- Fragnière, A. (2016). Climate change and individual duties. *WIREs Clim Change*, 7: 798-814. <http://doi:10.1002/wcc.422>
- Kingston, E., Sinnott-Armstrong, W. (2018). What's Wrong with Joyguzzling? *Ethical Theory and Moral Practice*, 21 (1): 169-186. <https://doi.org/10.1007/s10677-017-9859-1>
- Parfit, D. (1984). *Reasons and persons*, Oxford: Oxford University Press
- Sinnott-Armstrong, W. (2005). It's not my fault. Global warming and individual moral obligations". In (ed.), Sinnott-Armstrong W., Howarth R., *Perspectives on Climate Change*, Elsevier. Reprinted in (ed.), Gardiner S.M., Caney S., Jamieson D., Shue, H. (2010), *Climate Ethics – Essential Readings*: 332-346.