

Against Denialism

John Broome*

ABSTRACT

Several philosophers deny that an individual person's emissions of greenhouse gas do any harm; I call these "individual denialists." I argue that each individual's emissions may do harm, and that they certainly do expected harm. I respond to the denialists' arguments.

1. INTRODUCTION

Some people deny that humanity as a whole does harm by its emissions of greenhouse gas. I call these "species denialists." Others deny that individual human beings do harm by their emissions. I call these "individual denialists."

Individual denialism is surprisingly common among moral philosophers.¹ This paper opposes it.² I start in section 2 by reviewing some of the relevant science and economics. In particular I describe the significance of the atmosphere's extreme instability. In section 3, I draw what I believe to be the correct conclusion: an individual's emissions may or may not do harm, but they certainly increase the expectation of harm. The denialist's claim that they do no harm is not true in general.

I then go on to examine various arguments for individual denialism that can be found in the philosophical literature, and respond to them. I start in section 4 with a collection of very influential arguments that come from Walter Sinnott-Armstrong, and recently from him in collaboration with Ewan Kingston. In section 5, I answer an argument from overdetermination, which I attribute to Elizabeth Cripps. In section 6, I examine an argument from imperceptibility: that the harm done by an individual is not perceptible, and that it is therefore no harm at all because there are no imperceptible harms. I use an argument of Derek Parfit's to demonstrate there are indeed imperceptible harms, and I close one lacuna in his argument. An appendix to this paper criticizes Shelly Kagan's different, widely-cited response to the argument from imperceptibility.

2. SCIENCE AND ECONOMICS

Suppose that for fun you go for a ride in your gas-guzzling SUV on a Sunday afternoon.³ Say you use 10 litres of fuel, which means you emit around 25 kilos of carbon

*University of Oxford and Australian National University

dioxide. Economists have estimated what they call “the social cost of carbon” (SCC), which is a monetary measure of the harm done around the world by carbon dioxide emissions. The Obama administration’s figure for the SCC was about \$40 per tonne.⁴ It implies that the harm done by your drive is about \$1. Over your lifetime, if you are a typical academic who travels, you will probably emit more than 1,000 tonnes of carbon dioxide. At \$40 per tonne, the harm done by your lifetime emissions will probably be more than \$40,000. Your joyguzzling contributes a little to this sum.

The figure of \$40 per tonne is contentious. In this issue, Marc Fleurbaey et al. survey some of the problems with it. The very idea of setting a monetary value on many of the harms done by climate change is dubious. For example, it makes no sense to measure in money the harm that is done to wild animals. Many people also think that the preservation of nature has a value of its own, apart from the benefit it brings to human beings and animals, and this value cannot be measured in money. Some harms are more easily monetized, though Fleurbaey et al. explain that economists sometimes use dubious methods for monetizing them. The figure of \$40 per tonne probably gives a misleadingly benign impression of the harm done by carbon dioxide, because so much is left out of it. At any rate, the harm is not zero. Even if you dislike the whole idea of setting a monetary value on the harm, you should recognize that the harm done per tonne of carbon dioxide is not nothing.

Economists arrive at their figure by using models to estimate a “damage function,” which predicts the amount of global harm—the harm that will be done around the world and through all time—as a function of total emissions in each year. The SCC of emissions in a particular year is the first derivative of this global harm with respect to emissions in that year. It is the rate at which increasing emissions in that year increases global harm. Emissions in different years have different SCCs; the \$40 figure is for contemporary emissions.

The damage function from which it is derived is estimated from very large-scale models. This large-scale perspective overlooks the many irregularities that perturb the function on a small scale. The climate is complex, and so is the relation between emissions and harm at a small scale.

Concentrate on that Sunday afternoon, and imagine drawing a graph of global harm as a function of the total amount of carbon dioxide emitted on that one day. Hold constant emissions on all other days at what they actually are. Seen on the scale of millions of tonnes, this graph would be like [figure 1](#): a straight line with an upward slope equal to the SCC.

Do not read this graph and those that follow as showing how harm develops over time. The graphs are counterfactual rather than temporal. The vertical axis shows all harm around the globe, aggregated across all time. The horizontal dimension shows the various different amounts of carbon dioxide that might be emitted on one Sunday. The amounts emitted on other days are kept constant.

[Figure 1](#) shows a smooth line. But if it were magnified to show the effects of a few thousand tonnes or less it would be jagged. The SCC averages out the jaggedness.

Why would it be jagged? Two alternative accounts are available. The first is that climate change does some of its harm through bigger or smaller discrete events, such

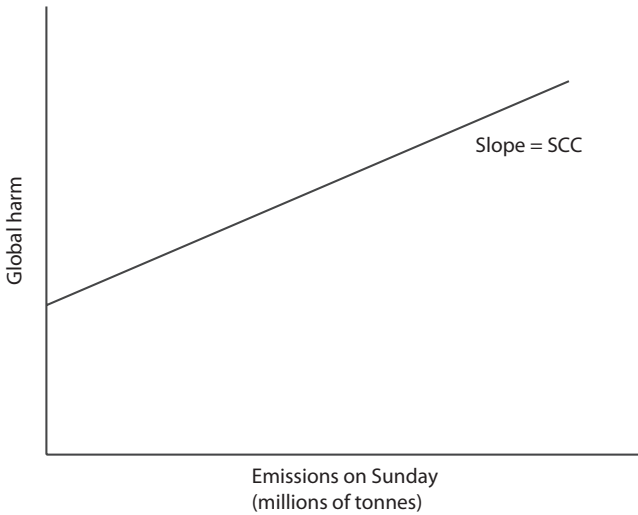


Figure 1. Damage function at a large scale

as a typhoon or a child's death from cholera. It does other harms through continuous processes. For instance, as water tables sink, people have to do more and more exhausting work to get the water they need; as the sea level rises, arable land is gradually washed away and families grow more hungry. This suggests that the graph of harm will slope steadily upwards, and be punctuated by upwards jumps. There may also be downward jumps. For instance, a small change in emissions may change the course of a storm, so that it does less harm. Figure 2 illustrates the graph for this account of jaggedness.⁵

That first account ignores the atmosphere's extreme instability. The atmosphere is a chaotic system, which means that a small disturbance at one time and place can escalate to cause very large disturbances in the weather all over the world just a few weeks later. In 1972 Edward Lorenz gave a lecture entitled "Predictability: Does the Flap of a Butterfly's Wings in Brazil Set Off a Tornado in Texas?" and proposed the answer, "it might."⁶ It remains an unresolved question in meteorology whether a disturbance as small as a butterfly-flap can really escalate to a global scale.⁷ If it cannot, the reason is that the disturbance of a butterfly-flap is on such a small scale that the viscosity of the air may damp out its effect. For the atmosphere, the scale on which viscosity is significant is less than a centimetre.

But the ten litres of fuel you burn during your Sunday drive contain about 340 million joules of energy, nearly all of which will be dissipated into the atmosphere in one way or another. Perhaps 50 million joules of it will be directly expended in stirring up the air. This is perhaps a trillion times the energy of a butterfly-flap. And that is the tip of the iceberg. Over the next century, the carbon dioxide emitted by your drive will cause more than a trillion joules of energy from the sun to be absorbed by the earth.⁸ Some of this will warm the atmosphere and continue to stir it up. So the doubts that arise over the butterfly effect do not extend to joyguzzling.⁹

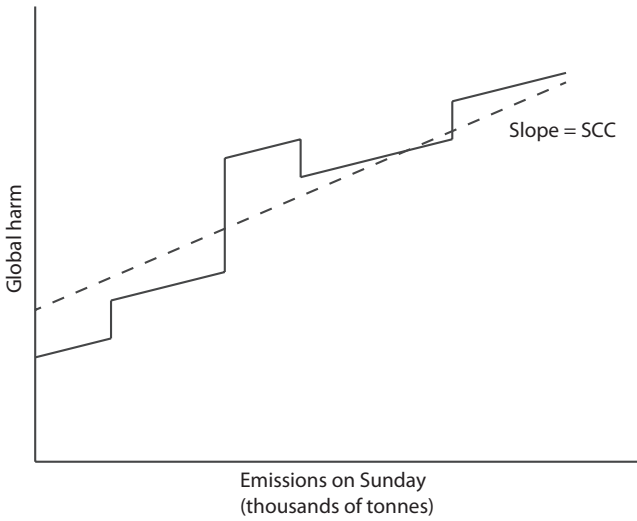


Figure 2. Damage function at a small scale, without chaos

Given the atmosphere's instability, we should expect global weather in a few decades' time to be entirely different if you go joyguzzling on Sunday from what it would have been had you stayed at home. Increasing emissions does not cause continuous changes punctuated by occasional discrete events such as a typhoon or a child's death from cholera, as the first account of jaggedness supposes. Instead it will cause typhoons to form at quite different times and places, and it will lead to a completely different distribution of cholera outbreaks. Your Sunday drive will cause a completely different group of people to be exposed to cholera and other risks of death. Some who would have died will survive because of your drive, and others who would have survived will die. The total numbers who die, and the total amount of harm done in the world may also be greatly altered.

A graph of the damage function will therefore oscillate chaotically up and down. Figure 3 is an illustration. This is the second account of its jaggedness. It will have no steady upward slope anywhere. There is literally zero probability that emitting 25 kilos will do no harm and no good. (This does not mean it cannot happen.¹⁰) Also, there is about equal probability that it will do good as that it will do harm.

The same models of the atmosphere as those that exhibit chaotic behaviour for the weather do not exhibit it for the climate. The climate consists of long-run averages of weather, and these are much more stable and predictable. It is a feature of chaos that the weather, following its unpredictable course, roughly follows stable cycles known as "attractors." Since the harm done by emissions of greenhouse gas is a total over centuries, it benefits from this stability of averages. So, although the graph of the damage function oscillates wildly up and down, the oscillations will not be huge. Even for small-scale changes, the graph will not deviate hugely from the straight, upward-sloping line that represents the average.

Most writing on the ethics of climate change ignores the instability of the atmosphere. In order to keep in touch with this writing, in this paper I shall consider the

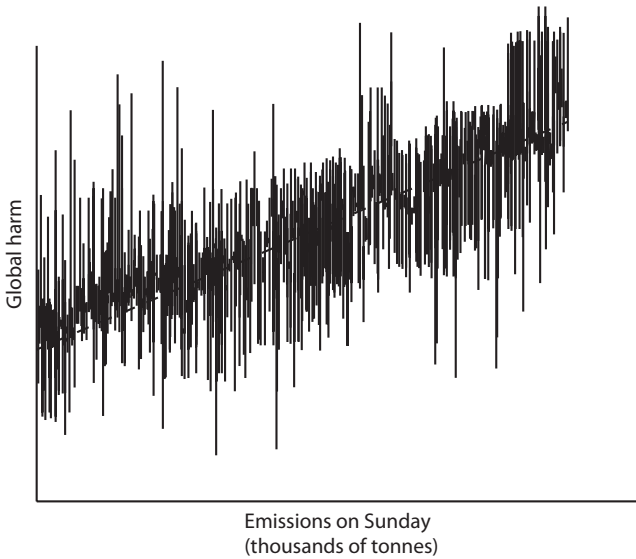


Figure 3. Damage function at a small scale, with chaos

consequences of both accounts of why the damage function is jagged: the one that takes account of instability and the one that does not. Nevertheless, given the atmosphere's chaotic nature, and given the decades that pass between the date of a particular emission of greenhouse gas and many of the harms that result from it, it is safe to assume that the chaotic picture of the damage function is the right one. The correct account of its jaggedness is the second.

3. EXPECTED HARM

Explained either way, the jaggedness implies that, when you consider whether or not to joyguzzle that Sunday afternoon, you cannot know what good or harm will actually result from what you do. The result may be a typhoon or a child's death, or it may be good. In the face of uncertainty like this, what you ought to do depends, not on the goodness of actual results, which you cannot know, but instead on the goodness of the "prospect" that each of your alternative acts leads to. A prospect is a portfolio of all the various outcomes that might result from an act, each associated with its probability of happening.

How prospects should be valued is the subject of expected utility theory.¹¹ This theory assigns to each possible outcome something called its "utility." "Utility" is a technical term defined within the theory; its meaning is described in the next paragraph. The theory demonstrates on the basis of some axioms that the value of prospects is given by their expectation of utility. That is to say: one prospect has more value than another if and only if it has a greater expectation of utility (or "expected utility"). "Expectation" has its standard mathematical meaning: the expectation of utility is a weighted average of the utilities of the prospect's possible outcomes, each weighted by its probability.

The utility of an outcome is not defined as its value, but it is a strictly increasing function of its value: higher value implies higher utility. Distinguishing utility from value allows the value of a prospect to depend on its degree of riskiness. Two prospects that have the same expectation of value may have different expected utilities. In particular, the more risky one, which includes a greater spread of values among its possible outcomes, may have a lower expected utility. This will be so if the graph of utility as a function of value, which slopes upwards because utility increases with value, has a downward curvature. In that case the more risky prospect is worse than the less risky one, even though they have the same expectation of value.¹²

The argument that the value of a prospect is given by expected utility does not depend on intuition. It is demonstrated in expected utility theory on the basis of axioms, so it rests on the truth of the axioms. The axioms can be defended,¹³ so the argument has a good foundation.

A case can be made for denying the distinction between utility and value, on the grounds that we have no basis for a quantitative scale of value apart from utility.¹⁴ In any case, the distinction makes no difference to the argument of this paper, so here I shall ignore it. I shall assume that the value of a prospect is given by its expected value. I use “good” as synonymous with “value,” and “harm” for negative value. The conclusion is that, where there is uncertainty, it is expected good and expected harm, rather than actual good and actual harm, that contribute to determining what you ought to do.

Your joyguzzling on Sunday afternoon creates a prospect that has a positive expectation of harm to other people. Its expectation of harm is given by the SCC, which measures the average, or expected, harm done by emissions of carbon dioxide.

Whether or not you *ought* to joyguzzle may depend on various things. No doubt joyguzzling brings some benefit to you, and this benefit may be worth more than the \$1 of expected harm it brings other people. On the other hand, it may be that you ought not to expose people to even a small expectation of harm just for your own enjoyment. Whatever the right conclusion about what you ought to do, one thing definitely counts against joyguzzling, and that is the expectation of harm it creates. Your act may or may not do harm, but it certainly creates an expectation of harm. Individual denialists to not claim merely that your emissions may not do harm, which is true. They claim they actually do no harm, which is not true in general.

4. SINNOTT-ARMSTRONG AND KINGSTON

Walter Sinnott-Armstrong’s “It’s not *my* fault” popularized individual denialism among moral philosophers. In this paper I have found the following arguments for it.

- Global warming will still occur even if I do not drive just for fun. Moreover, even if I do drive a gas-guzzler just for fun for a long time, global warming will not occur unless lots of other people also expel greenhouse gases. So my individual act is neither necessary nor sufficient for global warming (Sinnott-Armstrong 2005, 297).

Global warming will occur even if you do not joyguzzle, but there will be less of it. Your individual act is sufficient for increasing expected global warming. Similarly, there will be murders even if you refrain from murder, but there will be fewer of them. An individual murder is sufficient for increasing the number of murders.

- The harms of global warming result from the massive quantities of greenhouse gases in the atmosphere. Greenhouse gases (such as carbon dioxide and water vapor) are perfectly fine in small quantities. They help plants grow. The problem emerges only when there is too much of them. But my joyride by itself does not cause the massive quantities that are harmful (Sinnott-Armstrong 2005, 298).

It may be that when total emissions were small, increasing them a bit did no harm. The damage function may be flat or even downward-sloping for small emissions. But that was in the Eighteenth Century. Increasing emissions now do positive harm.

- Global warming is more like a river that is going to flood downstream because of torrential rains. I pour a quart of water into the river upstream My act of pouring the quart into the river is not a cause of the flood. Analogously, my act of driving for fun is not a cause of global warming (Sinnott-Armstrong 2005, 298–99).

Your quart could have been the straw that broke the camel's back, or more accurately the water that broke the dam and caused a flood. Sinnott-Armstrong simply stipulates that the flood will occur anyway. Even so, the more the water, the worse the flood. Your quart may make no difference to the harm done, but it may alternatively be just enough to cause a person to lose her grip and be washed away and drowned, or it may bring down another house, or do some other harm. In any case, global warming is not like a river that is going to flood; the atmosphere is more chaotic.

- You might think that my driving on Sunday raises the temperature of the globe by an infinitesimal amount. I doubt that, but, even if it does, my exhaust on that Sunday does not cause any climate change at all. No storms or floods or droughts or heat waves can be traced to my individual act of driving. It is these climate changes that cause harms to people. Global warming by itself causes no harm without climate change (Sinnott-Armstrong 2005, 299).

Having an effect is not the same as having an effect that can be traced.¹⁵ Your joyguzzling will have an effect but it will be untraceable. It may trigger storms or floods or droughts, but you can be confident that these disasters will not be traceable to you. That does not mean you do not cause them.

- There is nothing bad about global warming or climate change in itself if no people (or animals) are harmed. But there is no individual person or animal who will be

worse off if I drive than if I do not drive my gas-guzzler just for fun. Global warming and climate change occur on such a massive scale that my individual driving makes no difference to the welfare of anyone (Sinnott-Armstrong 2005, 301).

It is possible that no individual people or animals will be worse off, but this is very unlikely in view of the atmosphere's instability. Given the instability, some people will be hit by storms, droughts, or epidemics who would otherwise have emerged unscathed (and others will be saved from these disasters). Even setting instability aside, your joyguzzling will do some harm to people and animals through continuous processes such as the falling water table in some parts of the world. It may also trigger an upward jump in the damage function; it may cause a child to die from cholera, for example. In any case, it increases the expectation of harm.

In a new paper by Ewan Kingston and Walter Sinnott-Armstrong,¹⁶ I have found another two arguments.

- Our opponents have not shown that global warming is not [an] emergent [property] (Kingston and Sinnott-Armstrong 2018, 176).

These authors suggest that climate change may be like the sliminess of oil: individual molecules of oil are not slimy; sliminess emerges only when many molecules are put together. But physics shows that global warming is not like that. Each molecule of greenhouse gas absorbs those photons that encounter it, if their wavelengths are within a particular band. The more molecules there are, the more photons are absorbed. Global warming is simply the accumulated effect of absorption by individual molecules.

- 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 . . . 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. . . . 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. An injury from a burst seawall, or a death from expanded malaria range, are just as grievous if they happen at time t or $t + 1s$ (Kingston and Sinnott-Armstrong 2018, 177).

Provided concentrations of carbon dioxide in the atmosphere are increasing, it is true that the effect of cancelling a small act of emission is to delay slightly the time when each particular concentration is reached. It will also lower the peak concentration, if a peak is eventually reached.¹⁷ If we ignore the atmosphere's instability, the effect is to delay slightly the time of each harm that in due course will result from global warming, and if a peak is reached to prevent some harms altogether. All this

implies that cancelling the emission makes the total harm suffered at each time less than it would have been. The total benefit of cancelling the emission is the integral over all future times of the reduction in harm at each time. This integral is what is measured by the social cost of carbon. Italicizing “fraction of a second” changes nothing.

You might now wonder whether the benefit of cancelling the emission is infinite—quite contrary to Kingston and Sinnott-Armstrong’s claim that it is zero. Once carbon dioxide is in the atmosphere, some fraction of it remains there in effect forever. So cancelling an emission can be expected to reduce by a small amount the harm done by climate change at each time forever. Its total benefit is therefore the integral over an infinite time of a positive amount.

However, economists who estimate the SCC estimate this integral as finite, because they discount future benefits exponentially. Exponential discounting leads to a convergent integral. One justification they offer for discounting is the greater uncertainty of the further future: benefits predicted further in the future are less likely to materialize. Another is that they expect future people to be better off than earlier ones, and they take benefits to be less valuable if they come to better-off people. How far discounting is justified is not a subject for this paper.¹⁸ My point in this paper is only that the integral is not zero.

All of this discussion of Kingston and Sinnott-Armstrong’s second argument ignored the atmosphere’s instability. Given that the atmosphere is actually unstable, the effect of cancelling an emission is unpredictable in its nature, size, and timing. All that can be said is—as before—that it diminishes the expectation of harm, and this is what matters.

5. OVERDETERMINATION

Two further arguments for denialism are popular in the literature. This section and the next examine them in turn.

It is sometimes said that the harm caused by climate change is “overdetermined.” An example of overdetermination is this: suppose three people push a rock off a cliff and wreak havoc in the village below. But suppose any two of them could and would have pushed the rock off had the third not joined in. Harm is done by all three pushing together, but it seems untrue of each one of them that her individual act did harm. The havoc would have been wreaked even if she had not pushed.

Elizabeth Cripps claims that the harm done by climate change is overdetermined.¹⁹ As I understand her argument, Cripps ignores the atmosphere’s instability. Also, she implicitly ignores continuously harmful processes, such as the steadily increasing difficulty of getting water in some parts of the world. She apparently assumes that the graph of the damage function consists of long, level stretches, punctuated by jumps. [Figure 4](#) illustrates the damage function as she imagines it.

If this were so, then if total emissions happened to be in the middle of a level stretch, there would be overdetermination. If you were to increase your emissions in those circumstances, your doing so would do no harm. Emissions would still be on the level stretch of the graph. The same is true of each person: increasing her

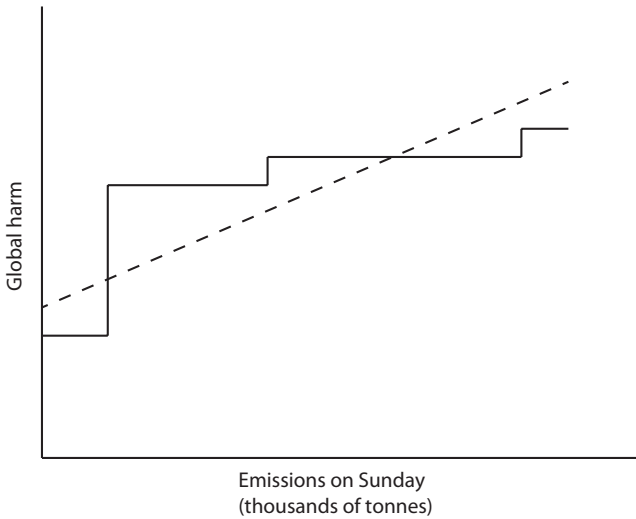


Figure 4. Damage function at a small scale without gradual harms

emissions would do no harm, even though a lot of harm would be done by all our emissions together. I assume this is Cripps's point.

However, if the total of everyone's emissions were not in the middle of a level stretch of the graph, but were instead just before a jump, then increasing your emissions would do harm because it would trigger the jump. If the total of emissions were just beyond a jump, reducing your emissions would do good.

Is it not very unlikely that total emissions are just before a jump? The answer may well be "yes" if we ignore the atmosphere's instability. Seen on the scale of an individual's emissions, jumps may be very spread out, so that the whole of a person's lifetime emissions are unlikely to trigger a jump. Does this not imply that you are very unlikely to do harm if you increase your emissions? Yes it does, provided we continue to ignore the atmosphere's instability and continue to ignore continuous harmful processes.

However, the fact that you are very unlikely to do harm—if it is a fact—does not by itself tell us how you ought to act. An elementary lesson from decision theory is that a decision should not be based only on what is very likely to result. Why should there be lifeboats on a ship? Not because they are very likely to be needed. They are very unlikely to be needed; most ships go to the breakers with their lifeboats still on board. The reason ships should carry lifeboats is that, in the very unlikely event of their being needed, the consequence of not having them would be dire. Similarly, other things being equal, you should avoid adding to your emissions because, in the perhaps very unlikely event that doing so will push total emissions over a jump in the damage function, the result would be dire. The jump may be the death of a child, or worse.

No one knows whether total emissions are just before a jump or not. When there is uncertainty, expected utility theory tell us we should act according to expectations. Increasing your emissions always has an expectation of harm, and this is true of each

person. Therefore—other things being equal—you should not do it. In terms of expected harm, there are no jumps and no overdetermination.

Against this, Cripps says:

Each individual is aware of others being motivated by their self-interested desires to perform the acts that would do harm in combination. Accordingly, the individual has every reason to believe that not only are there sufficient potential emitters, but that there are enough actual other emitters for her action to have only [a trivial impact]. Because the case is, in effect, overdetermined, each individual has reason to believe that, were all (or most) others to act in the relevant way, her actions would not trigger any extra harms. Because of the numbers involved and what she knows about others' motivations, she can assume that, were she to continue to emit at current levels, she would be one of just such a set.²⁰

Cripps appears to be arguing that, because you know other people are emitting just like you, you have reason to believe your emissions will not push the total of emissions past a jump. That is false. Other people's emissions add up to some quantity, and that quantity may be just short of a jump. This is so whether or not their emissions are like yours. You have no reason to believe it is not so. If the quantity is just short of a jump, your emissions will push the total past the jump.

However, Cripps may be getting at a different, more interesting point. Let us accept that, were you to act in a particular way, lots of other people would act that way. If you were to joyguzzle on Sunday afternoon, lots of other people would also joyguzzle that afternoon. This is not because you would cause them to joyguzzle. It is because they will joyguzzle anyway, and you would join their number.

Suppose the emissions of all those people and everyone else apart from you put the total of emissions just before a jump in the damage function. Then if you joyguzzle, the jump is passed. A child dies, say. Do you cause the child's death? Maybe not. It is true that the child dies and, had you not joyguzzled, the child would not have died. But this is also true of all the other people who joyguzzled that afternoon: the child dies and, had any one of them not joyguzzled, the child would not have died. Do each of you cause the child's death? Maybe not. The death is caused by all you joyguzzlers together, and maybe not by any one of you separately. If that is so, we might conclude that, even in the event that your action pushes total emissions past a jump, so harm results, you do not cause the harm. This may be Cripps's point. It is not really a matter of overdetermination but its opposite. It takes many people's actions together to determine the child's death.

Is this conclusion correct: is it true that, even if you push total emissions past a jump in the damage function, you do no harm? It is not. We are supposing that the total of other people's emissions is just short of a jump. This means that, were you to joyguzzle, a child would die and, were you not to joyguzzle, the child would not die. Whether or not you joyguzzle makes a difference to the world, and this difference is a child's death. If you joyguzzle, you make the world worse than it otherwise would have been, to this extent. The badness of the child's death therefore counts

against your joyguzzling in determining what you ought to do. If you know emissions are situated just before a jump, you ought not to joyguzzle unless this badness is counterbalanced by some stronger or equally strong consideration in favour of doing so.

Why is that? The obvious explanation is that your joyguzzling does harm to the extent of the child's death. Moreover, if you do not know whether your joyguzzling will push total emissions over a jump, this amount of harm goes into determining the expected harm of your joyguzzling, which is a basis for determining what you ought to do in that case. It is plain that joyguzzling does harm when it pushes the total of emissions past a jump.

We may have an intuitive conception of causation that makes us reluctant to say that you cause the child's death. If we said that, we would have to say the same about many other people, and we may be reluctant to accept that many people separately cause the same event. But this conception of causation is not morally relevant in this context. It should not debar us from accepting that your emissions do harm. The counterfactual truth that, were you to make the emissions, a child would die and, were you not to make them, the child would not die, is enough for that.

This counterfactual truth applies to everyone for whom it is the case that the acts of other people put the total of emissions just before a jump. If all of those people knew they were in this position, they would all have the same strong reason not to joyguzzle. Were any one of them to joyguzzle, a child would die, and were she not to joyguzzle, the child would not die. Each of these people are in this position because of the acts of other people including you if you joyguzzle. You are in this position because of the acts of other people including the joyguzzlers. In fact, all the potential joyguzzlers that Sunday afternoon are locked in a moral strategic game. To work out fully the moral duties they are under, we would need the tools of game theory.

But fortunately we do not need to employ those tools in the ethics of climate change. You are inevitably ignorant: you cannot know whether or not total emissions are just before a jump. Nor can anyone else. So you are not locked in a strategic game. All you know is the expected harm of joyguzzling. That is given by the social cost of carbon. This is what matters for whether or not you ought to joyguzzle. It is independent of the behaviour of other people that Sunday; it is very nearly constant. Because of this, you can correctly treat the rest of the world as a force of nature, not as a group of strategic agents interacting with you. In this respect, as in others, the ethics of climate change is comparatively straightforward. What counts against emitting, whenever and wherever you emit, is the constant social cost of carbon.

In any case, any argument based on overdetermination can get off the ground only if we assume away the atmosphere's instability and also assume climate change does none of its harm through continuous processes. These are incorrect assumptions. For that reason alone it is safe to ignore these arguments.

6. IMPERCEPTIBLE HARMS

I defined individual denialism as the claim that a person's emissions do no harm. Some philosophers espouse a softer form of denialism: they say the harm a person does by her emissions is insignificant.²¹ As a rough illustration of the harm you do, I

gave the figure of at least \$40,000 for an academic. This is not insignificant. But that was supposed to be total harm, and these authors are thinking of the harm you do to any particular person. They think this is insignificant. They draw the conclusion that the large figure of \$40,000, which is supposed to be the sum of all these amounts, must be mistaken.

Let us suppose, for the sake of argument, that the harm you do to each individual person around the world is small. Is it right to call it insignificant just because it is small? No. Some small amounts are very significant. The energy released by the fission of a single uranium atom is small, but the power of an atom bomb shows it is significant. It is significant because the small amounts of energy released by each of many atoms add up to something significant. We cannot tell whether a small amount is insignificant until we have checked whether many of these amounts add up to something significant. It is a fallacy to argue by first declaring the amounts insignificant and then drawing the conclusion on that basis that they cannot add up to something significant.

Since there are a lot of people in the world, even small harms done to each of them will add up to something significant. So a small harm cannot correctly be called insignificant. The softer form of denialism is therefore not available. If they are not to add up to something significant, harms will have to be not small but zero. You will have to do no harm at all to each person. Only zeros add up to zero. A denialist must say that your emissions do no harm at all.

It can be argued that the harm you do to any particular person through climate change is indeed zero, because it is so small as to be imperceptible. The argument is that an imperceptible harm is not a harm at all: a harm has to be perceptible. Consequently, since you do no perceptible harm to anyone, you do no harm in total. This is another defence of denialism. It depends on two claims: (1) that you do no perceptible harm through climate change, and (2) that there are no imperceptible harms. Both claims are false.

(1) is false because there are jumps in the damage function. Sometimes your emissions will trigger a jump, and many jumps are plainly perceptible. For instance, a jump might be the death of a child. This would be so even if the atmosphere were not unstable, and its instability multiplies the number of jumps. A small increase in emissions leads to completely different weather, and so to large increases and decreases in the harms that come to individuals. These are perceptible. (1) assumes that climate change is continuous in its effects, so the damage function is like [figure 5](#), but it is not. Notice, incidentally, that the damage function assumed by this defence of denialism is inconsistent with the one assumed by the defence from overdetermination, which is shown in [figure 4](#). The two defences are incompatible with each other.

Derek Parfit argues in *Reasons and Persons* that (2) is false.²² I think his conclusion is correct, but his argument contains a lacuna, which I shall repair in my own manner.

Here is Parfit's argument applied to a continuous effect of climate change. Suppose as a hypothesis for reductio that there are no imperceptible harms. Suppose the continuous fall of the water table at some place means that a person living there

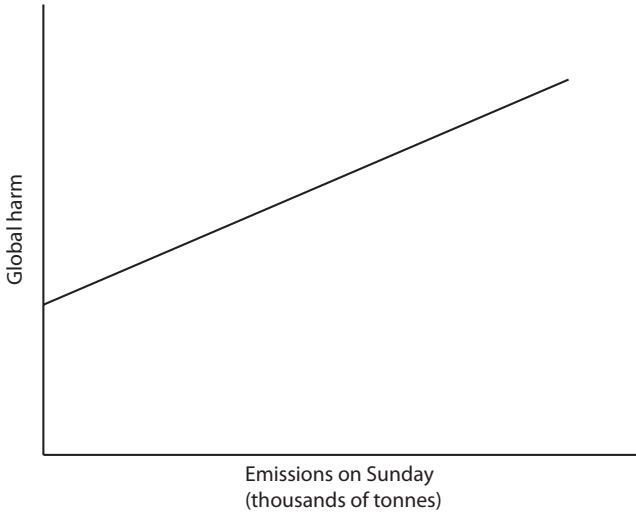


Figure 5. Damage function at a small scale, with only gradual harms

has to put in more and more painful effort each day to get water. Let d be a positive distance too small to be measured by any means. If the water table falls by this distance, the change must be imperceptible; were it perceptible, this perception would be a means of measuring the change. Suppose the water table drops from depth D to $D + d$. Since this change is imperceptible, it is harmless by hypothesis. The pain involved in getting water at depth $D + d$ is therefore not worse than the pain involved at depth D . Similarly, the pain involved at depth $D + 2d$ is not worse than the pain involved at depth $D + d$. Parfit assumes the not-worse-than relation is transitive,²³ which implies that the pain involved at depth $D + 2d$ is not worse than the pain involved at depth D . By iterating the argument, we can conclude that the pain involved at depth $D + nd$ is not worse than the pain involved at depth D , for any positive integer n . But for large enough n this is plainly false. So the hypothesis is false: there are some imperceptible harms.

In the example I have chosen, the conclusion of this argument is hard to believe. The harm in question is pain, and it is hard to believe that one pain is worse than another if the subject of the pain cannot feel any difference. In “Do I Make a Difference?” Shelly Kagan assumes this cannot be so.²⁴ His response is to deny that examples like mine are possible: he claims that at least one of the falls in the water table must be perceptible. I think his argument is mistaken, and I explain why in the appendix to this paper. I shall continue to assume that the example is possible.

Given the example, Parfit’s argument is valid. Its only lacuna is the premise that the not-worse-than relation is transitive. Many philosophers deny it.

Compare two pains of different qualities. One, A , is a deep ache; the other, B , an acute pang. Given some particular intensities for these pains, it is plausible that neither A nor B is worse than the other. Now suppose the acute pang becomes just a bit more acute, to make it $B+$. It is plausible that even $B+$ might not be worse than A . Pains are not so precisely comparable that a small worsening of one will necessarily

make it worse than the other. It follows that, though neither *A* nor *B* is worse than the other, they are not equally bad either. If *A* and *B* were equally bad, then *B*+, which is worse than *B*, would be worse than *A*, but it is not. Ruth Chang says that *A* and *B* are “on a par”;²⁵ I say they are “incommensurate.”

In this incommensurate situation, *B*+ is not worse than *A*, *A* is not worse than *B*, but *B*+ is worse than *B*. So the not-worse-than relation is not transitive.

In that example the intransitivity is generated by comparing pains that have different qualities: an ache and a pang. The water-table example does not involve pains of different qualities, so you might think a similar intransitivity could not be generated there. But it can.

Even if we think of pains with the same quality, such as aching muscles, it is plausible that the quantity of these pains is nebulous, so that it is not always sharply determined whether or not a particular pain is worse than another. Here is a crude model that represents this thought. The badness of a pain is represented by an interval of numbers rather than a single number; this is supposed to capture its nebulousness. For instance, one pain *A* might be represented by the interval $[9, 11[$ and another *B* by the interval $[8, 10[$. One pain is worse than another if and only if the whole of its interval is above the whole of the other’s interval. Then $B = [8, 10[$ is not worse than $A = [9, 11[$, which is not worse than $B+ = [10, 12[$, but $B = [8, 10[$ is worse than $B+ = [10, 12[$. The not-worse-than relation is again intransitive, but nothing requires these pains to have different qualities.

Examples like these lead many philosophers to reject Parfit’s assumption that the not-worse-than relation is transitive. Consequently they reject his argument against (2).²⁶ However, I think that the not-worse-than relation is indeed transitive. I think these apparent counterexamples are misdescribed. What is portrayed as incommensurateness is actually vagueness. My argument for this claim appears in my paper “Is Incommensurability Vagueness?”²⁷ The worseness relation is vague. In the examples, it is vague whether *B* is worse than *A*, and vague whether *A* is worse than *B*+. In the second example, there is a simple scale of badness but it is vague where on the scale of badness each pain lies.

I adopt a supervaluationist account of vagueness.²⁸ According to this account, a vague worseness relation can be thought of as an assemblage of putative worseness relations that are each sharp rather than vague. Each is a putative sharpening of the actual, vague worseness relation. A proposition about the relative badness of pains is actually true—true of the actual worseness relation—if and only if it is true according to all the sharpenings. In particular, it is true that a pain *A* is worse than another *B* if and only if *A* is worse than *B* according to all the sharpenings. It is true that *A* is not worse than *B*—which is to say it is false that *A* is worse than *B*—if and only if *A* is not worse than *B* according to all the sharpenings. If *A* is worse than *B* according to some sharpenings, and not worse than *B* according to other sharpenings, then it is neither true nor false that *A* is worse than *B*.

Because I am offering vagueness as an alternative to incommensurateness, I make an assumption that rules out incommensurateness. I assume that each of the sharpenings is a complete ordering.²⁹ That is to say, in any of the sharpenings, of any two pains *A* and *B*, either *A* is worse than *B*, or *B* is worse than *A*, or else *A* and *B* are

equally bad. There is no incommensurateness in any the sharpenings. This is an extra assumption that is not part of supervalueationism itself. It implies that there is no incommensurateness in the actual worseness relation, since nothing is true of the actual relation unless it is true of all the sharpenings. Supervalueationism together with this assumption of completeness abolishes incommensurateness entirely.

For instance, in the first example as I first described it, I said that *A* is not worse than *B*, *B* is not worse than *A*, and *A* and *B* are not equally bad. That was to say that *A* and *B* are incommensurate. As I now describe the example, it is neither true nor false that *A* is worse than *B*, and neither true nor false that *B* is worse than *A*.

Because every sharpening of the worseness relation is a complete ordering, in every sharpening the not-worse-than relation is transitive. Suppose that in some sharpening *X* is not worse than *Y* and *Y* is not worse than *Z*. Then either *Y* is worse than *X* or *X* and *Y* are equally bad, and either *Z* is worse than *Y* or *Y* and *Z* are equally bad. In all cases, either *Z* is worse than *X* or *X* and *Z* are equally bad. So *X* is not worse than *Z* in this sharpening.

Since the not-worse-than relation is transitive in every sharpening, supervalueationism implies it is transitive in the actual, vague worseness relation. Replacing incommensurateness with vagueness therefore vindicates Parfit's premise. I have drawn this conclusion on the basis of the supervalueationist account of vagueness, but I believe it could also be drawn on the epistemicist account too.³⁰ So if I am right that what is often taken to be incommensurateness is actually vagueness, Parfit's demonstration that there are imperceptible harms is sound. His conclusion is true even though it is hard to believe.

I recognize that this argument for the transitivity of the not-worse-than relation is contentious, and it relies on assumptions that I have not justified in this paper. So it is worth mentioning that a doubt about transitivity arises only when the badness in question is nebulous, like pain. The harms of climate change are not all like this—even the continuous ones. For example, the harm done by a continuous process might be to shorten a person's life in imperceptible steps. The not-worse-than relation in this case is equivalent to the not-shorter-than relation applied to the length of the person's life. This relation is uncontentiously transitive because the length of a life is precise, not nebulous. Parfit's argument applied to this example shows that imperceptible shortenings of life are harmful. This should be uncontentious, and it is enough to show that claim (2) is false: there are some imperceptible harms.

The defence of denialism on grounds of imperceptibility depends on claims (1) and (2). I have now concluded that both are false. Since the argument against (2) is not entirely uncontentious, I put most weight on the falsity of (1). This defence can apply only to the continuous processes of harm, but even if we set aside the atmosphere's instability, many of the harms of climate change are not continuous. They are discrete events such as the death of a child, and these are plainly perceptible. This defence of denialism is therefore unsuccessful.

7. CONCLUSION

Individual denialism is damaging. It may be widely believed among the public: many nonphilosophers may believe that their own emissions do no harm. If so, it may contribute to the public's apathetic response to climate change.

Individual denialism among moral philosophers is damaging in a different way. Not many philosophical denialists are apathetic. Many of them think that each of us should reduce our emissions. But they find it hard to explain why we should, given that they think our emissions are harmless.

Their denialism leads them to think that the ethics of climate change is more complicated than it is. They explore notions of group responsibility, group action, complicity and so on, searching for a reason to reduce emissions. These are fascinating topics for a philosopher. They are also important and worth exploring in their own right. But there is a straightforward reason for us to reduce our emissions: they do expected harm. So these fascinating topics are beside the point.

Denialism has distracted some of the moral philosophers who work on climate change; it has persuaded them to expend energy on such irrelevant topics. They would do better to turn their attention to topics within the philosophical theory of value that are genuinely important for the subject. Some of those topics are illustrated by the papers in this volume.³¹

NOTES

1. For example, denialist arguments appear in Cripps (2013, 119–24); Kingston and Sinnott-Armstrong (2018); Maltais (2013); and Sinnott-Armstrong (2005).
2. It joins the opposition expressed in this journal by Avram Hiller (Hiller 2011).
3. I take this famous example from Sinnott-Armstrong (2005).
4. www.obamawhitehouse.archives.gov/blog/2013/11/01/refining-estimates-social-cost-carbon
5. This is the account of jaggedness I gave in Broome (2016).
6. The lecture is based on Lorenz (1963), and more particularly Lorenz (1969).
7. See T.N. Palmer et al. 2014.
8. Figure 8.29 in the report of IPCC Working Group 1, p. 712, shows an “absolute global warming potential” for carbon dioxide, calculated over 100 years as 30×10^{-13} watt-years per square metre per kilo. I calculate on that basis.
9. See T.N. Palmer et al. 2014, section 5.
10. If you spin a perfectly balanced pointer, there is zero probability that it will come to a stop pointing exactly north. But it could happen.
11. Expected utility theory is presented by R.A. Briggs (2014). I interpret it as a theory of value rather than a theory of choice; this interpretation is explained in my *Weighing Goods* (1991).
12. In this volume, Lara Buchak (2018) describes a different way of understanding the value of avoiding risk.
13. See Broome 1991.
14. See Broome 1991, chapter 10; Greaves 2015.
15. As Garrett Cullity (2018) also says in this volume.
16. Kingston and Walter Sinnott-Armstrong 2018.
17. As Aaron Maltais (2013) recognizes when he makes a similar point.
18. See Fleurbaey et al. 2018, section 3.1.
19. Cripps 2013, 123.
20. Cripps 2013, 123–24. Cripps gives credit to Aaron Maltais (2013).
21. For instance, Joakim Sanberg (2011).
22. Parfit 1984, 79.
23. Parfit 1984, 78.
24. And so does Julia Nefsky (2012) in her response to Kagan.

25. Chang 2001.
26. For instance, Spiekermann (2014).
27. See Broome 1998.
28. See Fine 1975.
29. This assumption is defended in Broome (1998).
30. See Williamson 1994.
31. My thanks to Renée Bolinger, Elizabeth Cripps, Garrett Cullity, Caspar Hare, Tom Hurka, Douglas MacLean, Matthew Rendall, and Tim Palmer for helpful advice. I particularly thank Sæde Hormio, whose thesis *Marginal Participation, Complicity, and Agnotology* (2017) has been a very useful guide to me, and who has also given me extensive valuable advice on this paper. Research for this paper was supported by ARC Discovery Grants DP140102468 and DP180100355.

REFERENCES

- Briggs, R.A. 2014. "Normative Theories of Rational Choice: Expected Utility," *Stanford Encyclopedia of Philosophy*.
- Broome, John 1991. *Weighing Goods: Equality, Uncertainty and Time*, Blackwell.
- . 1998. "Is Incommensurability Vagueness?" in *Incommensurability, Incomparability, and Practical Reason*, ed. Ruth Chang, Cambridge, MA: Harvard University Press, 67–89.
- . 2016. "A Reply to My Critics," *Midwest Studies in Philosophy* 40: 158–71.
- Buchak, Lara 2018. "Weighing the Risks of Climate Change," *The Monist* 102: 66–83.
- Chang, Ruth 2001. "The Possibility of Parity," *Ethics* 112: 659–88.
- Cripps, Elizabeth 2013. *Climate Change and the Moral Agent: Individual Duties in an Interdependent World*, Oxford University Press.
- Cullity, Garrett 2018. "Climate Harms," *The Monist* 102: 22–41.
- Fine, Kit 1975. "Vagueness, Truth and Logic," *Synthese* 30: 265–300.
- Fleurbaey, Marc, et al. 2018. "The Social Cost of Carbon: Valuing Inequality, Risk, and Population for Climate Policy," *The Monist* 102: 84–109.
- Greaves, Hilary 2015. "Antiprioritarianism," *Utilitas* 27: 1–42.
- Hiller, Avram 2011. "Climate Change and Individual Responsibility," *The Monist* 94: 349–68.
- Hormio, Sæde 2017. *Marginal Participation, Complicity, and Agnotology: What Climate Change Can Teach Us About Individual and Collective Responsibility*, PhD Thesis, University of Helsinki.
- Intergovernmental Panel on Climate Change, Working Group 1, 2013. *Climate Change 2013: The Physical Science Basis*, Cambridge University Press.
- Kagan, Shelly 2011. "Do I Make a Difference?" *Philosophy and Public Affairs* 39: 105–41.
- Kingston, Ewan, and Walter Sinnott-Armstrong 2018. "What's Wrong with Joyguzzling?" *Ethical Theory and Moral Practice* 21: 169–86.
- Lorenz, Edward 1963. "Deterministic Nonperiodic Flow," *Journal of the Atmospheric Sciences* 20: 130–41.
- . 1969. "The Predictability of a Flow Which Possesses Many Scales of Motion," *Tellus* 3: 290–307.
- . 1993. "Predictability: Does the Flap of a Butterfly's Wings in Brazil Set Off a Tornado in Texas?" in *Essence of Chaos*, University of Washington Press, Appendix 1 (orig. 1972).
- Maltais, Aaron 2013. "Radically Non-Ideal Climate Politics and the Obligation to at Least Vote Green," *Environmental Values* 22: 589–608.
- Nefsky, Julia 2012. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan," *Philosophy and Public Affairs* 39: 364–95.
- Palmer, T.N., A. Döring, and G. Seregin 2014. "The Real Butterfly Effect," *Nonlinearity* 27: R123–R141.
- Parfit, Derek 1984. *Reasons and Persons*, Oxford University Press.

- Sandberg, Joakim 2011. "My Emissions Make No Difference: Climate Change and the Argument from Inconsequentialism," *Environmental Ethics* 33: 229–48.
- Sinnott-Armstrong, Walter 2005. "It's Not My Fault: Global Warming and Individual Moral Obligations," in *Perspectives on Climate Change: Science, Economics, Politics, Ethics. Advances in the Economics of Environmental Research*, Volume 5, ed. W. Sinnott-Armstrong and R.B. Howarth, Elsevier, 293–315.
- Spiekermann, Kai 2014. "Small Impacts and Imperceptible Effects: Causing Harm with Others," *Midwest Studies in Philosophy* 38.
- Williamson, Timothy 1994. *Vagueness*, Routledge.

APPENDIX: KAGAN AND IMPERCEPTIBILITY

In the example in section 6, I assumed that the fall, d , in the water table from one day to the next is small enough to be imperceptible, yet the larger fall nd over n days is perceptible. In "Do I Make a Difference?" (Kagan 2011, sections 11–15) Shelly Kagan denies this is possible. Given that the fall over n days is perceptible, he claims that, for at least one day, the fall between that day and the next must be perceptible.

Actually, I made the stronger assumption that d is immeasurable by any means. If a fall d is immeasurable in this sense, it is imperceptible, because perception is a means of measurement. Kagan is therefore committed to denying that my stronger assumption is possible. I shall start with this stronger assumption because it gives Kagan's argument the best chance. I shall return to imperceptibility at the end of this appendix.

Here is Kagan's argument, transposed to measurement. Suppose the depth of the water table is measured by the best laser device, which displays the depth in metres to ten decimal places. Suppose that over n days, the display changes from 10 metres to 11 metres. The argument proceeds by reductio. On day 1, the display reads "10.0000000000." On day 2 the water table has dropped a distance d , which is by hypothesis immeasurable, so the display still shows the same reading. The hypothesis implies that it also shows the same reading on day 3, and so on to day n . But on day n it shows a different reading, so the hypothesis is false.

There must be at least one day when the reading changes. Indeed, there must be many such days, but let us concentrate on the first. There must be a first day—say day r —when the device will show a different reading from "10.0000000000," namely "10.0000000001." The reading on day r is different from the reading on day $r-1$. Yet the difference in the depth of the water table between those two days is d . So the device has measured a distance as small as d .

The conclusion of this argument is that no positive distance is small enough to be immeasurable. That is not credible. But the argument looks strong. What is wrong with it?

As it approaches a borderline, in practice a digital display starts to flick from one reading to another. We might imagine an idealized device that behaves as I have described, and moves from an unwavering "10.0000000000" to an unwavering "10.0000000001" from one day to the next. But if d is small enough, this is to imagine that the universe is deterministic, whereas actually it is not. For very fine

measurement, the flicking of the display results from quantum variations in the process of measurement. The idealized device is impossible.

The difference between the device's display on one day and its display on the next is a difference in frequency. On a later day, the device reads "10.0000000001" rather than "10.0000000000" more frequently than on an earlier day. When the depth of the water table gets above 10.0000000005, the frequency of "10.0000000001" will rise above 50% if the device is as good as possible. But if d is small enough, it will be undetectable which day this happens on, because of statistical variations. The difference in frequency between one day and the next will be undetectable in the time available, which is to say a day. For a very small d it will be undetectable in the history of the universe. This is why a small d is immeasurable.

What about the first day when "10.0000000001" appears with any positive frequency at all? That day is surely distinguishable from the previous one. But actually that is day 1. From the very beginning there is a small positive probability that the device will display this reading; that is a feature of quantum mechanics. This number will flash up for some very short time even on day 1.

Come back now to imperceptibility. The fall in the water table causes increasing pain. Kagan claims that the increase in pain between some one day the next must be perceptible. His argument is the same. The subject's response to the pain on day n differs from her response on day 1. Therefore there must be one day when it differs from the day before. This shows the subject must perceive a difference in pain between those two days.

The failure of this argument is the same. The subject's response may be indeterminate. The only change between one day and the next may be that the subject has a slightly higher propensity to complain. This difference of propensity may be undetectable because of statistical variations. All this is parallel to the case of measurement.

But in this case the indeterminism is psychological and need not arise from quantum indeterminism. For instance, the subject's brain state might vary according to deterministic variations in the magnetic field.