

Time for Caution

Johanna Thoma

October 4, 2021

I

Policy-makers ought to make laws aiming to phase out or restrict activities that could plausibly lead to catastrophe. Individuals ought not engage in activities that pose a significant threat to the lives of others. Prudentially, they should not make choices that come with a substantial risk of financial ruin. Precautionary principles such as these are frequently appealed to both in policy-making and in discussions of good individual decision-making. They prescribe omission or reduction of an activity, or taking precautionary measures, whenever potential harmful effects of the activity surpass some threshold of likelihood and severity.

In the policy context, precautionary principles can come to clash with what is often regarded as the gold standard of policy evaluation, namely risk-cost-benefit analysis. In risk-cost-benefit-analysis, for each policy option, we add up all potential costs and benefits, weighted by the probability with which we take them to occur. In contrast, when applying a precautionary principle, once we have determined that a potential harm stemming from an activity is serious and likely enough, the potential benefits that may come from taking the risk are disregarded. The activity should be avoided, restricted, or precautions should be taken, even if the potential benefits of the activity or the costs of precaution are substantial. According to proponents of precautionary principles, this is just as it should be: They take risk-cost-benefit analysis to be insufficiently precautionary.

*I am grateful to Richard Bradley, the members of the LSE Philosophy Department's Early Career Work-in-Progress Seminar, and participants at the Workshop in Practical Philosophy at Saarland University for very helpful comments on earlier drafts of this paper. I also benefitted greatly from discussions with Dominic Alford-Duguid and Lavinia Picollo, and thank audiences at the University of Bern, the University of Freiburg, the University of St Andrews, and the LSE Choice Group for very helpful questions and feedback.

A parallel observation applies to precautionary principles in individual decision-making and the orthodox theory of rational choice in the context of risk and uncertainty, expected utility theory. Under the standard interpretation, expected utility theory asks agents to take into account their evaluation of each of the potential outcomes of a proposed act by weighting it by its probability. In contrast, when we apply a precautionary principle, once we have determined that some particularly harmful potential outcomes are likely enough, it no longer matters how good all other potential outcomes are: The action should not be performed, or at least not without taking potentially costly precautions. To those who are sceptical of expected utility theory as a theory of individually rational choice for being insufficiently permissive of risk averse attitudes, this could be an advantage. In the moral context, deontologists, in particular, are often sceptical of attempts to incorporate expected utility theory when determining our moral obligations under risk. Precautionary principles appear to be promising ways of expressing deontological restrictions in the context of risk and uncertainty.¹

Given the centrality of precautionary principles in public policy discussion (in particular in environmental policy and the regulation of new technologies)² and their appeal in

¹Indeed, as I explore more in the following, there are several important insights in debates on deontology under risk in the ethics literature that are relevant for the discussion of precaution in the policy and individual context. Sergio Tenenbaum, for instance, argues against what he calls the ‘multiplicative model’ for choice under uncertainty for deontologists, where the multiplicative model weights evaluations of the different consequences of actions by their probabilities. Sergio Tenenbaum. Action, deontology, and risk: Against the multiplicative model. *Ethics*, 127(3):674–707, 2017. Instead, he proposes that deontological restrictions apply to act-types described in a way that incorporates their riskiness. For instance, he suggests ‘One may not endanger someone’s life except in order to secure a very significant good’ as a candidate deontological restriction in the context of uncertainty, where ‘endangering’ must be understood as posing more than a trivial risk to somebody else’s life (p.697). This principle has the structure of a precautionary principle. Others have been more critical, e.g. Frank Jackson and Michael Smith who argue against such ‘threshold absolutism’, while acknowledging it might be the best way to approach risk and uncertainty from a deontological perspective. Frank Jackson and Michael Smith. Absolutist moral theories and uncertainty. *The Journal of Philosophy*, 103(6):267–283, 2006.

²Discussion of precautionary principles arose in the context of German environmental policy in the 1970s, under the name of the *Vorsorgeprinzip*. Internationally, precautionary principles rose to prominence when they were discussed at the International Conference on the Protection of the North Sea in 1984 and 1987, and influenced the Ministerial Declarations resulting from the conferences. In 1992, two important events enshrined precautionary principles as guiding principles for environmental policy both internationally, and in many nations world-wide. First, the EU made reference to ‘the’ precautionary principle in its founding document, the Maastricht Treaty. Second, the *Rio Declaration* following the United Nations Conference on Environment and Development in Rio de Janeiro advocated the “precautionary approach” to the management of forests. Precautionary principles have since been appealed to in contexts as diverse as climate change mitigation, (a precautionary principle is referred to in the latest IPCC (2014) AR5 WGII report), approval of medical trials, regulation of chemicals and new technologies, preemptive warfare (George W. Bush’s justification of the invasion of Iraq, which pointed to the possibility of weapons of mass destruction in the country’s possession, is often thought to be an instance of precautionary reasoning – see Craig McLean, Alan Patterson, and John Williams. Risk assessment, policy-making and the limits of knowledge: The precautionary principle and international relations. *International Relations*, 23(4):548–566, 2009), and protection of animal welfare (see Jonathan Birch. Animal sentience and the precautionary principle. *Animal Sentience*, 16(1):1–15, 2017).

the individual realm, precautionary principles deserve serious consideration as principles of choice in the context of risk and uncertainty, constraining or supplanting the application of risk-cost-benefit analysis and expected utility theory. I will be concerned with three desiderata for precautionary principles in the individual and policy-making context: First, precautionary principles should not make grossly implausible recommendations. Second, precautionary principles should provide useful action-guidance to agents, that is, it should be easy for agents to work out what the principle requires them to do or not do. And third, precautionary principles should be *effective* action-guiding principles, by which I mean that agents who have accepted the normative authority of a precautionary principle and worked out what it recommends should not find it overly difficult to go through with its recommendations. The first desideratum is essential. The second and third need not be, depending on the purpose for which one is proposing the precautionary principle – after all, there are many moral principles that are imperfectly action-guiding, and expected utility theory itself is arguably too complex to implement to be effectively action-guiding. If the only purpose of precautionary principles was to express true moral principles within a deontological framework, or to express true principles of rationality that capture greater risk aversion, it may be desirable, but not important that these principles are themselves effectively action-guiding. Their implementation can potentially be dealt with separately, through action-guiding rules of thumb.

The second and third desiderata are crucial, however, for what I take to be the core purpose of precautionary principles in the public policy context, namely, to be effective decision-making and legal tools to guard against procrastinating on or altogether failing to confront and mitigate certain kinds of risk. In the policy realm, precautionary principles are most frequently invoked to guard against risks that are hard to quantify – because the harm is very unlikely, because the scientific evidence on the possibility and likelihood of harm is yet slim, or because the potential harms stem from radical changes to human and non-human life that we find it hard to evaluate – and whose quantification there is consequently little agreement on. These risks may be very important to consider for a policy-maker. Yet, not being able to precisely quantify the potential harms and benefits of an activity and their likelihoods is a serious obstacle to conducting traditional risk-cost-benefit analysis. And lack of agreement on correct quantification will decrease the public credibility of any attempted evaluation. While there may be ways to extend the traditional framework to incorporate imprecise values and probabilities, there is a temptation in practice to simply ignore those risks we find it hard to quantify. This temptation will be even greater when there is reason to think that the state of our knowledge regarding the risk will improve over time, making a ‘wait and see’ strategy attractive. Proponents of precautionary principles worry about this appeal to lack of scientific knowledge and quantifiability serving as an excuse for either dangerous policy procrastination, or a reck-

less inattention to those risks whose quantification is difficult or controversial.³ In the individual context, too, there can be a temptation to use the fact that certain risks are hard to evaluate precisely as an excuse for delay, inaction or recklessness.

The advantage of precautionary principles, by contrast, is that they seem not to be as vulnerable to this excuse. All they apparently require for their application is a determination of whether the potential harms from an activity pass some, usually vaguely defined, threshold of severity and likelihood. This determination is often much less controversial than any particular risk-cost-benefit or expected utility analysis would be. For instance, while there is an overwhelming consensus in the relevant scientific community that there is a plausible mechanism by which current levels of global CO2 emissions can lead to catastrophe if maintained, there is little agreement on the precise social cost of carbon. If determining the former is sufficient to yield a clear prescription to severely restrict activities that produce CO2 emissions, then there is much less room for excuses to delay action than if any action needed to be backed up by a risk-cost-benefit analysis. For this reason, in fact even those who take some variant of risk-cost-benefit analysis or expected utility theory to be the true measure of the value of policy options might support some precautionary principles as action-guiding principles for pragmatic purposes in certain contexts.⁴ But precautionary principles will only serve as effective tools against policy procrastination, inaction or recklessness in the face of risk if they are effectively action-guiding.

In the following, I explore an important obstacle for precautionary principles fulfilling the role of effective and plausible action-guiding principles in the face of risk and

³In fact, one famous early statement of the ‘precautionary approach’ to public policy, namely that found in the 1992 Rio Declaration, defined the approach as ruling out using scientific uncertainty as a reason to postpone precautionary measures. In a recent policy brief, the European Commission describes its commitment to the precautionary principle as a tool for decision-making and legal instrument. They write, “incomplete information, inconclusive evidence and public controversy can make it difficult to achieve consensus over the appropriate response to hazardous substances or activities, but these are precisely the sorts of conditions that often demand hard and fast decisions.” (p.3) And, “the principle can be invoked in cases when regulators have to take decisions in advance of scientific ‘certainty’ on an issue or risk, or to create the impetus to take a decision by removing excuses for inaction on the grounds of scientific knowns.” (p.20) European Commission (Environment Directorate General) and University of the West of England (Science Communication Unit). The precautionary principle: Decision-making under uncertainty. *Science for Environmental Policy: Future Briefs*, 18, 2017. Lyda Akinyi Omuko describes how in environmental litigation, the precautionary principle can assist courts in stopping defendants from exploiting scientific uncertainty as an excuse. Lydia Akinyi Omuko. Applying the precautionary principle to address the “proof problem” in climate change litigation. *Tilburg Law Review*, 52, 2016. In the philosophical literature concerned with policy applications, too, the precautionary principle is often defended as a practical tool for overcoming obstacles to action. Dan Steel, for instance, defends it as a corrective in the light of a historical record in environmental policy of more frequent ultimately costly failures to take preventative action as opposed to costly rushes to preventative action. Daniel Steel. *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy*. Cambridge University Press, 2014.

⁴For a recent defence of this idea, see Andreas Christiansen. Rationality, expected utility theory and the precautionary principle. *Ethics, Policy and Environment*, 22(1):3–20, 2019.

uncertainty. My starting point here is the following observation: Many of the risks that precautionary principles are meant to guard against are not the consequence of only one single choice. On the policy scale, if, or sadly rather when the consequences of catastrophic climate change occur, this is and will be the result of decades of repeated failure to take preventative action by policy-makers, corporations and individuals around the globe. If a country's drinking water becomes polluted enough to pose serious health risks, this is likely the result of various different contaminants released into the water through various different pathways, and the responsible policy-maker's failure to adequately regulate each of them. If the commercial or political use of advanced neuro-technology came to undermine our basic rights to freedom of thought and privacy, this would be the result of an extended process of research and development, approval and implementation of many different technologies, which could have been restricted at various points in time and in various different ways. If a country fails to contain the spread of a deadly virus, this is likely the result of multiple failures of policy and logistics at many levels and over an extended period of time. On the individual scale, serious risks to one's own and other people's health and lives are often the result of repeatedly engaging in dangerous or unhealthy activities.

Given the fact that risks that are sufficiently harmful and sufficiently likely to fulfil the antecedent of a precautionary principle often accumulate over time, those advocating precautionary principles need to explain how the principles are meant to effectively govern decision-making over time. I will argue for two claims in the following. First, to yield the normative verdicts proponents of precautionary principles would like to make, precautionary principles must be understood to be diachronic principles, which requires some added structure to how they are commonly formulated. And secondly, such diachronic precautionary principles still face serious obstacles in their implementation. In the best cases, carefully formulated precautionary principles might be useful as normative background principles that, in the policy context, can enjoy wide public support. But they are less likely, on their own, to be effective tools to counter (policy) inaction, procrastination or recklessness in the face of risk. They still need to be supplemented with other decision-making tools. Moreover, there is one class of cases, where incremental risks are continuously resolved before new risks are taken, where it is doubtful whether precautionary principles have any purchase at all.

II

The most frequently cited appeals to a precautionary principle, such as the 1998 Wingspread Declaration, or the Maastricht Treaty, refer to 'the' precautionary principle in the singular, and intend it to apply to a wide variety of both individual and policy

choice contexts where harms to human health or the environment are at stake. However, it has also been acknowledged that more specific formulations of the principle are needed for a precautionary principle to offer substantial action-guidance to individuals and policy-makers. And there seems to be a broad agreement amongst proponents of precautionary principles today that different kinds of precautionary principle are appropriate for different contexts. In particular, proponents of precaution typically accept that the urgency of precaution, and the kind of precaution required will depend on the severity of potential harm. Where the survival of mankind is at stake, less evidence of risk should be required to take precautions than when we are ‘merely’ facing a regional, non-lethal health hazard. Lauren Hartzell-Nichols has consequently proposed that we should no longer speak of “the” precautionary principle, but rather of precautionary principles in the plural, all of which are instances of a “precautionary approach”, or precautionary reasoning.⁵ I will follow this use here.

What matters for our purposes is that most precautionary principles proposed in the literature share some common structure, and are meant to express the spirit of precaution in the face of uncertain harms – of being risk averse, and being ‘rather safe than sorry’. The shared structural features of precautionary principles that I am concerned with here are the following:

1. There is a threshold of likelihood and/or severity of potential harm caused by one’s activities beyond which precautionary measures (e.g. omission) are required.
2. This threshold is/these thresholds are vague. That is, they allow for borderline cases that do not clearly fall above or below the threshold.
3. The principle does not specify explicitly the scope of the activity it is meant to be applied to. In particular, it is not made explicit whether the principle is meant to be applied to small-scale individual choices, or also to entire courses of action over time, and if the latter, to how long and comprehensive a course of action.

Consider the following three examples of precautionary principles:

- a. The option with the best worst-case outcome that meets some minimal level of plausibility should be chosen.⁶
- b. “If a scientifically plausible mechanism exists whereby an activity can lead to a catastrophe, then that activity should be phased out or significantly restricted.”⁷

⁵Lauren Hartzell-Nichols. From ‘the’ precautionary principle to precautionary principles”. *Ethics, Policy and Environment*, 16(3):308–320, 2013.

⁶Sven-Ove Hansson. The limits of precaution. *Foundations of Science*, 2:293–306, 1997.

⁷Daniel Steel. *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy*. Cambridge University Press, 2014, p.28.

- c. “One may not endanger someone’s life except in order to secure a very significant good.”⁸

Each of these principles clearly shares features (1), (2) and (3) above. (a) is a maximin principle, but one that only takes into account risks that pass a vague threshold of likelihood. (b) refers to a vague threshold of harm, namely ‘catastrophe’, that needs to be met for us to be required to phase out or significantly restrict an activity that might cause it. Moreover, ‘scientifically plausible mechanism’ is at least in part a threshold of likelihood. (a) and (b) are meant to be applied to ‘options’ and ‘activities’ respectively, leaving the scope of the activity unspecified. (c) rules out ‘endangering’ somebody else’s life, unless a very significant good is at stake. And Sergio Tenenbaum, who puts forward (c), proposes that an act only counts as ‘endangering’ once it poses more than a trivial risk to somebody else’s life, that is, once it passes a vague threshold of likelihood of harm, which may also be sensitive to context. The required precaution once this threshold is crossed is omission of the activity. The principle itself does not make reference to the scope of the activity it is meant to apply to.⁹

There are also good reasons why these principles share at least features (1) and (2). Vague thresholds of likelihood such as the one in (a) are often defended as an application of a ‘de minimis’ condition, which is a concept that entered more traditional risk-cost-benefit analysis in the 1980s. In that time, the FDA started explicitly treating some risks, such as a one-in-a-million lifetime risk of death from exposure to a substance, as ‘negligibly small’, and excluding them from analysis.¹⁰ As Sven-Ove Hansson argues, it makes sense for precautionary principles to similarly require that risks pass at least a *de minimis* likelihood threshold before precaution can be required, for two reasons.¹¹ First, the strict worst-case consequences of all options available to us are arguably equally catastrophic:

⁸Sergio Tenenbaum. Action, deontology, and risk: Against the multiplicative model. *Ethics*, 127(3):674–707, 2017, p.697.

⁹As we will see below, Tenenbaum in fact intends it to apply to any scope of activity.

¹⁰For an early defence, see Cyril L. Comar. Risk: A pragmatic *de minimis* approach. *Science*, 203(4378):319, 1979.

¹¹A *de minimis* condition for risks to warrant precaution has been defended by various other proponents of precautionary principles. Some critics of the precautionary principle have attacked a naive ‘catastrophe principle’, which proposes taking precautions against any event or activity that has any possibility of leading to (personal or collective) catastrophe. Clearly, such a principle would be extreme in requiring us to take precautions against all sorts of remote possibilities. As Christian Munthe notes, it would be absurd to take precautions against a 1-in-19 billion lifetime risk of developing cancer from a food additive. Christian Munthe. *The Price of Precaution and the Ethics of Risk*. Springer Netherlands, 2011. Moreover, the catastrophe principle has been argued to be incoherent, since any precautionary measure itself has the possibility of leading to catastrophe. This was argued most prominently by Cass Sunstein. C. Sunstein. *Laws of Fear: Beyond the Precautionary Principle*. Cambridge University Press, 2005. Amongst others, Martin Petersen and Katie Steele respond to these worries by defending a *de minimis* clause. Martin Peterson. What is de minimis risk? *Risk Management*, 4(2):47–55, 2002 and Katie Steele. The precautionary principle: A new approach to public decision-making? *Law, Probability and Risk*, 5(1):19–31, 2006.

Logically speaking, any decision may have unforeseen catastrophic consequences. If far-reaching indirect effects are taken into account, then – given the unpredictable nature of actual causation – any decision may lead to a nuclear holocaust. Any action whatsoever might invoke the wrath of evil spirits (that might exist), thus drawing misfortune upon all of us.¹²

And, second, even if we did have options that completely rule out such extremely unlikely worst-case scenarios, this may detract from precautions that are intuitively more important. For instance, in the policy context, if we judge that the immediate annihilation of all life on Earth through collision with a large asteroid is worse than the worst-case consequences of global climate change, then a precautionary principle without a *de minimis* clause may recommend spending a massive amount of resources on asteroid impact avoidance, which could otherwise be used for climate change mitigation and adaptation.

While some formulations of precautionary principles recommend precaution regarding all risks that pass a *de minimis* threshold (so are non-negligible), others feature higher likelihood thresholds than non-negligibility for special precautionary measures to be recommended. As (b) and (c) exemplify, many precautionary principles have what Daniel Steel calls a ‘tripod structure’.¹³ This tripod structure specifies a harm condition (a degree of seriousness of harm), a knowledge condition (a state of knowledge we have about the potential harm occurring), and a precautionary measure that is proposed if the harm condition and the knowledge condition are met. This structure then allows us to express what appropriate precaution amounts to in a particular context, potentially allowing us to flexibly address a broader variety of policy and individual choice contexts. For instance, (b) is proposed by Steel as plausibly applicable to climate change mitigation. In this case, the harm condition is that there is a risk of catastrophe, the knowledge condition is that there is a scientifically plausible mechanism by which an activity leads to the harm, and the proposed remedy is that the activity be phased out or significantly restricted. Clearly, all precautionary principles that share this tripod structure will share feature (1) above, since they specify a harm and knowledge condition after which precaution is recommended.

Insofar as they are precisely formulated, a common worry about *de minimis* rules and other thresholds of likelihood and harm is that they are arbitrary.¹⁴ There seems to be no important difference between a one-in-a-million lifetime risk of death from a substance and a one-in-a-million-and-one risk. Yet, strictly speaking, this would make all the difference for FDA regulation. In the deontology under risk literature, the arbitrariness

¹²Sven-Ove Hansson. The limits of precaution. *Foundations of Science*, 2:293–306, 1997, p.300.

¹³Daniel Steel. *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy*. Cambridge University Press, 2014. See also, e.g., Christian Munthe. *The Price of Precaution and the Ethics of Risk*. Springer Netherlands, 2011.

¹⁴See, e.g., Kristin Shrader-Frechette. Technological risk and small probabilities. *Journal of Business Ethics*, 4:431–445, 1985.

of thresholds in deontological principles that take the form of precautionary principles is also often criticised.¹⁵ In response to the arbitrariness worry regarding *de minimis* clauses, Martin Peterson argues that we should acknowledge that the *de minimis* condition is in fact vague, and will thus have borderline cases.¹⁶ As in other cases of vagueness, this does not mean that there aren't clear cases of risks that are or are not *de minimis*.

A more important reason for the vagueness of thresholds is that characterisations of thresholds of likelihood and harm in precautionary principles often refer to qualitative concepts that simply are vague. For instance, Hansson proposes “scientific reasonableness” as a precautionary threshold.¹⁷ Tenenbaum aims to rule out “endangering” lives.¹⁸ Examples of alternative harm conditions are that consequences are ‘serious’, ‘harmful to humans’, ‘irreversible’, or ‘such as to reduce or eliminate biodiversity’.¹⁹ There are evidently borderline cases of ‘scientific reasonableness’, ‘endangering’, ‘serious’ harm, or of ‘catastrophe’. Yet, these vague concepts may be the normatively most relevant categories for a particular context. Even when we put some effort into defining these conditions more precisely, it will be difficult to eliminate vagueness altogether. For instance, Hartzell-Nichols defines ‘catastrophic’ outcomes as those “in which many millions of people could suffer severely harmful outcomes.”²⁰ This is helpful in understanding what a catastrophe is, but clearly both “many millions of people” and “severely harmful outcomes” admit of borderline cases themselves. That they feature vague thresholds of likelihood and/or harm, and therefore fulfil conditions (1) and (2) above, thus seems essential to precautionary principles as they are most commonly discussed.

Finally, it helps to remember that precautionary principles are often invoked in cases of deep uncertainty, where we cannot assign precise probabilities to all outcomes, perhaps because we lack conclusive scientific evidence, or cannot precisely evaluate the potential harms involved. If our judgements of likelihood or of harm are imprecise,²¹ then even if we were to formulate a precise likelihood or harm threshold, there will be borderline cases. This is because for some risks, it will not be clear whether they fall above or below the harm or likelihood threshold. The presence of such borderline cases, even if thresholds were precise, is enough to trigger the problems stemming from vagueness that I will discuss

¹⁵E.g., by Frank Jackson and Michael Smith. Absolutist moral theories and uncertainty. *The Journal of Philosophy*, 103(6):267–283, 2006.

¹⁶Martin Peterson. What is *de minimis* risk? *Risk Management*, 4(2):47–55, 2002.

¹⁷Sven-Ove Hansson. The limits of precaution. *Foundations of Science*, 2:293–306, 1997, p.300.

¹⁸Sergio Tenenbaum. Action, deontology, and risk: Against the multiplicative model. *Ethics*, 127(3):674–707, 2017.

¹⁹See Neil A. Manson. Formulating the precautionary principle. *Environmental Ethics*, 24:263–274, 2002.

²⁰Lauren Hartzell-Nichols. Precaution and solar radiation management. *Ethics, Policy and Environment*, 15(2):158–171, 2012, p.160.

²¹Imprecision in credences is often represented with a family of probability functions. This can lead to borderline cases of the application of a precautionary principle when some members of this family assign a probability above the threshold, and others below the threshold.

in the following.

III

I have shown that typically, precautionary principles specify a vague threshold of harm and/or likelihood beyond which precaution is required. The core challenge I want to raise here is that this makes it possible that we find ourselves in situations where many different actions together result in a risk of harm severe enough to trigger the application of a particular precautionary principle, while no individual action seems to do so. Since, as we just saw, precautionary principles tend not to specify the scope of activity they are meant to apply to, this results in two problems. First, which scope or scopes of activity should the precautionary principle be applied to? I will refer to this as the ‘problem of scope’. And second, can precautionary principles effectively govern an agent’s decision-making over time? I will refer to this as the ‘problem of execution’. To illustrate, consider the following three stylised examples, the first two of individual decision-making, and the last of policy-making.

Investing in Your Friends’ Startups. Your 100 friends want you to invest in their 100 startups. You judge that for each, the chance of failure is non-negligible. But you wouldn’t lose much money on each, and the expected payoff from each investment is positive. Moreover, you like to help out your friends. And so much speaks in favour of investing in each. You also think that the probability of failure of the different startups is independent, and that it is very likely you would make an overall gain from investing in all. However, you know that it is still possible that all of the startups fail. This would mean financial ruin for you. You are a cautious person, and care deeply about avoiding financial ruin. You formulate the following precautionary principle:

PP-Investment: If there is a plausible chance that some choice leads to financial ruin, then I should not make that choice.

You judge that the antecedent is true for all 100 investments considered together: There is a plausible chance that all fail, which would lead to a clear case of financial ruin. At the same time, you judge that the antecedent is never true for any individual investment in a startup. The potential loss from each investment is small, and you judge that no individual added investment will ever turn the worst-case scenario for your overall portfolio from non-financial ruin into financial ruin.

Practising for a Piano Exam. You’ve been learning the piano for some years, and you are starting to think that you want to pursue a career as a

professional pianist. A minimal condition for becoming a professional pianist is that you pass an upcoming performance exam. Passing the exam is thus crucial for your future life design. You are confident that you will pass it with sufficient practice, and that sufficient practice wouldn't be so disruptive as to undermine any other important life goals. But you also know that if you don't practise enough, there is a substantial chance that you will fail the exam. You judge that generally speaking, practising more raises the chance of you passing your exam. But each time slot that you could use for practice could also be used for more immediately pleasurable activities, say, playing a computer game. You see a danger that you might fail to practise sufficiently for your exam, spending lots of time gaming instead. As a cautious person, you formulate this precautionary principle:

PP-Piano: If an activity with no lasting importance to me leads to a substantial chance that I will fail my piano exam, I should not engage in it.

Gaming has no lasting importance to you. And spending all available practice time slots before your exam gaming instead of practising would certainly lead to a substantial chance that you will fail the exam. And so the antecedent of this precautionary principle would hold true when considering the activity of gaming in all practice slots. However, it does not seem to hold true when considering gaming instead of practising on any individual occasion where you have an hour of available practice time: You judge that any one-hour practice session only has an insignificant and hard to identify effect on your chances of passing. In your estimation, no missed one-hour practice session takes you from a non-substantial to a substantial chance of failing.

Risky Technology. A large country currently relies for its basic functioning on a technology that is known to carry a risk of severe accident, resulting in the potential loss of life of dozens, and severe potential health consequences for thousands of people. There is an independent risk of this happening within the country of about 0.005% in any one month period of use, or about 0.00016% per day. Over a 10 year period, this would result in a risk of accident of about 0.6%. The risk could be removed immediately by stopping the use of the technology, and its use can be resumed at any time. As a policy-maker, you care very much about avoiding accidents. At the same time, the alternatives to relying on this technology are very costly, though still financially sustainable. You formulate the following precautionary principle, which you find to have broad public support:

PP-Technology: If an activity leads to a non-negligible risk of severe accident, and there is a significantly safer and financially sustainable alternative available, we should pursue the alternative.

Suppose that, were it not for the risk of accident, continued use of the technology is the clearly favoured course of action. So the important question is whether the antecedent of your precautionary principle holds true for the activity of relying on the technology, that is, whether we judge the risk of accident to be negligible or not. Your problem now is the following: You judge that the risk of accident on any individual day (ballpark 0.00016%) or in any individual month (ballpark 0.005%) is negligible. However, the risk of accident over a 10 year period (ballpark 0.6%) is certainly not negligible.

Each of these cases features what I will call an ‘incremental activity’, namely investing in one startup, missing one hour of practice, and relying on the technology for one more month respectively. They also feature an extended activity which consists in repeatedly engaging in the incremental activity. What makes these cases problematic is that in each case, the proposed precautionary principle constrains the agent, by prescribing omission,²² when applied to the extended activity, but doesn’t constrain any of the incremental activities considered in isolation. At least this is so if we grant the judgement that no individual incremental activity can make the agent cross the harm or likelihood thresholds specified in the antecedent of the precautionary principle, an assumption I will call ‘tolerance’. Risks of the seriousness and likelihood that the precautionary principle aims to guard against accumulate over time in such a way that only extended courses of action seem to be constrained by the precautionary principle.

Call any such case a ‘cumulative risk case’. But note that the accumulation of the serious risk works differently in each of these cumulative risk cases. In the first case, it is the severity of potential harm that accumulates so that the harm threshold is only crossed when many incremental activities are considered together. Let us call such cases ‘cumulative potential harm cases’.²³ In the second and third cases, it is the likelihood of harm that accumulates, so that the likelihood threshold of the precautionary principle is only crossed when many incremental activities are considered together. Let us call such cases ‘cumulative likelihood cases’. Moreover, these two cases are distinguished by the manner in which the likelihood of harm accumulates. In the risky technology case, the community is exposed to risk every day, and this risk is continuously resolved. By the

²²The following discussion of the problem of scope will apply equally for other kinds of prescribed precaution (e.g. taking out insurance, wearing a helmet). As noted below, in the case of the problem of effective execution, some precautionary measures are not quite as prone to the problem as omission.

²³Note both the similarity to and the distinction from the much discussed ‘collective harm cases’, sometimes also referred to as ‘cumulative harm cases’. For an overview, see Julia Nefsky. Collective harm and the inefficacy problem. *Philosophy Compass*, 14:1–17, 2019. In collective harm cases many actions carried out by many different people together cause a severe harm *for certain*, but no individual choice seems to make a difference. Here many choices by the same agent together make it the case that the harms in a sufficiently likely worst-case scenario are severe enough to cross the vague harm threshold of a precautionary principle, but no individual choice seems to do so.

time the policy-maker decides whether to use the technology for another month, she will already know whether last month's use resulted in an accident or not. The risks from past activities are in that sense bygone – the community is not facing them anymore. By contrast, in the piano case, the likelihood of the harm of failing your exam accumulates without resolution, and is only resolved at the end of a long series of choices. To distinguish these two types of cases, call risky technology-type cases 'transient cumulative likelihood cases' and piano-type cases 'aggregating cumulative likelihood cases'.

All three types of cases also seem to occur in (real) policy contexts, though most real applications will of course be more complex than our stylised examples. For instance, the approval of clinical trials, decisions about road safety, or the continued operation of nuclear power stations seem to exhibit a similar pattern to the risky technology case and could be analysed as transient cumulative likelihood cases, where the risks of incremental activities, or of the continuation of a policy, are continuously resolved. Research into technologies with potential dual use, continuation of the activities that raise the probability of a runaway greenhouse effect, continuation of activities that provoke a hostile state, or loosening of restrictions during a global pandemic might involve a pattern similar to the piano exam case. As research results, emissions, memory of provocations, and opportunities for close contacts aggregate, at least for the relevant foreseeable future, these cases can involve an ever increasing likelihood of a severely harmful outcome at a later point in time. They could thus be analysed as 'aggregating cumulative likelihood cases'.²⁴ Lastly, climate change mitigation more generally also appears to involve a similar pattern to the investment case, where the worst-case scenarios involve extreme harms only because of many small-scale emissions decisions. To that extent, it can be analysed as a cumulative potential harm case.

As noted above, the judgement that the precautionary principles in our cases only restrict agents when applied to extended courses of action, and not when applied to incremental activities individually relies on a tolerance assumption: that incremental activities can't make an agent cross the harm or likelihood thresholds of a precautionary principle. I take this tolerance assumption to be especially plausible in the case of transient cumulative likelihood cases, as here tolerance can be justified independently of the vagueness of the thresholds. This is because in transient cumulative likelihood cases, the risk already incurred through past choices is bygone, already resolved. This makes it very plausible

²⁴However, real-life examples like this will also be more complex than the piano case in the sense that there is no one fixed point in time at which the dice will be thrown, as it were. Rather, once enough research, greenhouse gases, provocations, and opportunities for close contact have amassed to make the harm in principle plausible, there is some probability of research being used for harmful purposes, a runaway greenhouse effect being triggered, armed conflict breaking out, or exponential growth of infections at any point in time, and that probability increases through the continuation of the activity, and stays increased for some time. To the extent that a harmful effect could arise any day, and when it doesn't, that risk is bygone, such cases also exhibit the characteristics of transient cumulative likelihood cases.

that it is irrelevant for the evaluation of the next risky choice. The risk incurred in the next choice is not in any meaningful sense adding to some growing likelihood of harm (as is the case in the aggregating cumulative likelihood cases). If no accident has happened so far, and past risk has been resolved, the next risky choice only brings about a negligible risk for the agent, which supports the tolerance assumption.

In the other two cases, tolerance is plausible because of the vagueness of thresholds. In these cases, potential harm and likelihood aggregate without resolution. When we are yet far from reaching a threshold of harm or likelihood, another incremental activity won't make an agent cross it, and when it has already been crossed, another incremental activity is not what causes the substantial risk, and is thus not ruled out by the precautionary principle. However, in the borderline region, tolerance only seems plausible because, in the face of a vague threshold, incremental activities seem to make too insignificant a difference to, e.g., turn potential worst-case losses from non-financial ruin to financial ruin. If the tolerance assumptions in the above cases do not seem intuitively plausible, we can always individuate the incremental activities even more finely, e.g. considering only a fraction of an investment, one minute of missed practice time, or another day or minute of using the technology. That tolerance assumptions for small enough incremental changes seem intuitively very plausible is indeed a characteristic feature of vague predicates – this is what makes the Sorites Paradox a paradox.

Insofar as the efficacy of precautionary principles as action-guiding principles is concerned, the intuitive implausibility of denying tolerance is arguably already enough to make us doubtful that precautionary principles would *effectively* constrain individuals when applied only to incremental activities: The effectiveness of the principle would rely on agents identifying some specific incremental activity that brings them, e.g., from a case of potential of only non-ruinous financial harms to potential financial ruin, which seems too much to ask in practice. Nevertheless, motivated by aiming to resolve the Sorites Paradox, some theories of vagueness deny the truth of tolerance assumptions for vague predicates, in which case it might seem that, at least in theory precautionary principles may constrain incremental activities. In response to this concern, the Appendix argues that at least the three major theories of vagueness either don't support the idea that precautionary principles constrain some incremental activities considered individually, or would only do so under some strong additional assumptions.

I will thus take the intuitive analysis of our cases for granted, and moreover assume that a similar analysis applies in real world cases that share a similar structure, such as the policy applications considered above. For instance, a precautionary principle aiming to constrain activities that bring about a non-negligible chance of a catastrophic accident would only constrain the operation of nuclear power stations over long stretches of time, and not the activity of operating them for one more month or day considered in

isolation. A precautionary principle aiming to constrain activities that bring about a plausible chance of human extinction in the next few hundred years would only constrain CO2 emitting activities on a global scale considered over many years, but would not constrain incremental emitting activities considered in isolation. We can now formulate the problems of scope and execution more precisely.

The problem of scope arises because in the kinds of examples we considered, the scope of activity to which we apply the precautionary principle makes a difference. When we treat each individual investment, hour of practice missed or month of technology use as a separate activity and only apply the precautionary principle to those incremental activities, we come to a different conclusion about what the agent should do than when we bundle many missed practice hours or investments together in a larger choice problem, and apply the precautionary principle to this extended course of action. If we do the former, the agent is unconstrained by the precautionary principle, and free to engage in each incremental activity (which we assumed the agent would indeed do if unconstrained). If we do the latter, the agent is constrained not to engage in each incremental activity. Fixing the level(s) of scope at which the principle applies in different ways leads to divergent recommendations.²⁵ If that is so, how should precautionary principles be formulated and applied? Moreover, turning to the problem of execution, can precautionary principles be effectively action-guiding for extended courses of action?

IV

When, as we have seen, the scope or scopes of activities at which precautionary principles are applied makes a difference to what they recommend, how should the principles be formulated and applied? In the following, I will present some potential responses to the problem of scope, and eventually argue for the last one. If precautionary principles are to effectively guard against serious risks, they must be applied to extended series of choices over time in the kinds of examples we considered. But to do so in a way that has no implausible implications, some additional structure is needed to the way in which precautionary principles are formulated. I will argue that making the scope of activity to which the principle is meant to be applied explicit is most conducive to effective action-guidance. However, the following section will show that two more persistent obstacles to

²⁵Such scope dependence of recommendations also characterises alternative decision theories that may come apart from expected utility theory by allowing substantial risk aversion. Interestingly, in the case of rank-dependent expected utility theories such as Lara Buchak's risk-weighted expected utility theory, the divergence often goes the other way in similar cases to the ones considered here. Lara Buchak. *Risk and Rationality*. Oxford University Press, 2013. The theory may recommend an agent be risk averse for some small-scale decision considered in isolation, but recommend against a compound of many such decisions considered together. For discussion, see Johanna Thoma. Risk aversion and the long run. *Ethics*, 129(2):230–253, 2019.

effective action-guidance remain, one of which in fact calls into question whether diachronic precautionary principles can be rationally followed at all in some core cases.

A first potential response is to deny the problem. The last section already addressed one potential reason to deny the problem of scope exists, namely denying the tolerance assumption underlying the judgement that incremental activities are unconstrained when considered in isolation. Another potential strategy for denying the problem starts from the observation that most proponents of precautionary principles argue that different principles are appropriate for different contexts. The problem of scope could potentially be avoided if we successfully argue that different precautionary principles are appropriate when considering incremental activities or extended activities respectively. What we would need to show, in order to rule out the problem of scope (as well as the problem of execution), is that there is always a match between what the appropriate precautionary principle for the extended activity demands, and what the precautionary principle for the incremental activities demands. However, I take this to be highly implausible in our examples.

Usually, differing precautionary principles for different contexts are justified by an appeal to proportionality: what evidentiary standards apply should be calibrated to what precautionary measures we are demanding and to the severity of the potential harm in question.²⁶ Harm condition, knowledge condition and remedy should be chosen so as to express an appropriately precautionary attitude for the circumstances in question. In our risky technology example, the harm we are concerned with is the same whether considering incremental activities or the longer term activity: We want to avoid a serious accident. However, the proposed remedy for the incremental activity is less costly: It involves refraining from relying on the technology just for the next month, rather than a longer period. And so perhaps the knowledge condition can be less stringent for the incremental activity, so that the remedy is demanded even for a much less probable risk of accident – as low as the very low risk from using the technology for one month.

The problem with this response is that it implies that appropriate precaution would demand not engaging in the incremental activity even in circumstances where we face a truly one-off choice of whether to engage in it or not. But that seems implausible in many cases: In our example, the risk in any one month really does seem negligible in the context of a large country, and deciding not to rely on the technology, on a one-off occasion, seems overly cautious. Indeed, it is only when we start considering the consequences of long-term reliance on the technology that serious concerns about the risks of accident arise.²⁷ The same holds for the investment and piano practice decisions. A real-world

²⁶See, for instance, Daniel Steel. *Philosophy and the Precautionary Principle: Science, Evidence, and Environmental Policy*. Cambridge University Press, 2014, and Kerry H. Whiteside. *Precautionary Politics: Principle and Practice in Confronting Environmental Risk*. MIT Press, 2006.

²⁷Also see Sergio Tenenbaum. Action, deontology, and risk: Against the multiplicative model. *Ethics*,

application similar to the risky technology case where such judgements are commonly expressed is the debate around the potential dangers from radiation from backscatter x-ray scanners used in airport security.²⁸ The dosage of radiation from these scanners is low (much lower than the radiation from a medical x-ray scan), only to the surface of the body, and exposure is very short. But in theory, the radiation from each occasion of a backscatter x-ray scan could cause cell damage which could lead to cancer.²⁹ As in our risky technology case, the most common models of the cancer risks from radiation assume that each backscatter x-ray scan exposes passengers to an independent risk of developing cancer.³⁰ The EU banned backscatter x-ray scanners at airports in 2012, which may be interpreted as an application of the precautionary approach the EU is committed to. However, the justification of the ban was that at that point, a superior technology was available, which only exposed passenger to non-iodizing radiation – suggesting that the ban would not have been implemented if the costs to precaution had been higher. When it comes to advice to individual passengers, even the most concerned medical professionals only recommend refusing x-ray scans, and opting for a physical pat-down instead, to those especially at risk, and, importantly for us, frequent fliers.³¹ As in the risky technology case, their reasoning is that a very low risk only becomes a concern when we take it repeatedly, even though the risk on each occasion is independent from the risk on the next.

What does this mean for the response that, if we use the appropriate precautionary principle for each context, we will not get conflicting recommendations in the cases we are concerned with here? In these cases, refusing to take the one-off risk is not the proportionate precautionary reaction to the risk faced on one occasion. If precaution demands refusing the risk it is only because the agent faces the risk repeatedly. And so, if the appropriate precautionary principle for the context of the incremental decision of whether, for instance, to continue using the technology is such that it recommends not using it, it is the diachronic context of our past and future choices that makes it so. However, arguing that the diachronic context makes a difference in this way seems to be just another way of giving priority to precautionary principles applied to extended courses of action. And this is a way of having addressed the problem of scope, not a way of denying its existence.

Another reaction to the problem of scope might be that any agent is simply free to

127(3):674–707, 2017, p.13 for this analysis of a similar case involving cycling.

²⁸I thank [redacted] for this example.

²⁹Julie Accardo and M. Ahmad Chaudhry. Radiation exposure and privacy concerns surrounding full-body scanners at airports. *Journal of Radiation Research and Applied Sciences*, 7(2):198–200, 2014.

³⁰Pratik Mehta and Rebecca Smith-Bindmann. Airport full body screening: What is the risk? *Archives of Internal Medicine*, 171(12):1112–1115, 2011.

³¹David J. Brenner. Are x-ray backscatter scanners safe for airport passenger screening? for most individuals, probably yes, but a billion scans per year raises long-term public health concerns. *Radiology*, 259(1):6–10, 2011.

frame her decision problems as she sees fit, and should only apply the relevant precautionary principle to the decision as she framed it. For instance, in the investment case, you could either consider investing in each startup as a separate investment, or you could consider yourself as facing just one investment decision of how many and which startups to invest in. If you consider yourself as facing just one large investment decision, then *PP-Investment* will tell you not to invest in all startups, since that would bring with it a plausible chance of financial ruin. Instead, you should invest in few enough of the startups to rule out the possibility of financial ruin. What if you consider investing in each startup as a separate investment decision? Applied to each of those incremental activities separately, the precautionary principle does not recommend against investing. You are then free to invest in each, as we said you otherwise prefer.

There is in fact evidence that people's investment decisions can differ strongly depending on whether they are prompted to think of more long term or more short term returns, and so this frame dependence actually captures how many of us make decisions under risk and uncertainty.³² However, we typically think of this kind of frame dependence as problematic. In the kinds of contexts we are considering, if precautionary principles were only ever applied to the frame the agent chose, framing would make all the difference for how much precaution is actually recommended by the relevant precautionary principle. In fact, an agent prone to narrowly framing her decisions would not see her actions constrained very much at all by precautionary principles. But then those principles would not be doing their intended job as effective tools against (policy) procrastination, inaction, or recklessness. Moreover, it is counterintuitive that what an agent morally or prudentially ought to do should depend on a framing decision – unless, that is, there is a matter of fact about which frame an agent morally or prudentially ought to choose. But what the right scope of activity is to apply precautionary principles to is just the original problem of scope.

Another reaction to the problem of scope might be that in the kinds of examples we considered, precautionary principles lead to a problematic inconsistency in recommendation that ultimately counts against precautionary principles. In our cases, it might look like the agent ought to do incremental activity x_1 , and she ought to do incremental activity x_2 , and x_3, \dots but she ought not do (x_1 and x_2 and x_3, \dots). For instance, in the investment case, it might seem like you ought to invest in each of your friends' startups, but also that you ought not invest in all of them. Proposed deontological rules for choice under uncertainty are often criticised for such alleged failures of agglomeration.³³ However, this worry

³²See Shlomo Benartzi and Richard Thaler. Myopic loss aversion and the equity premium puzzle. *Management Science*, 45(3):364–381, 1999.

³³See, in particular, Frank Jackson and Michael Smith. Absolutist moral theories and uncertainty. *The Journal of Philosophy*, 103(6):267–283, 2006. My response here mirrors that of Patrick Hawley to Jackson and Smith. Patrick Hawley. Threshold absolutism defended. *The Journal of Philosophy*, 105(5):273–275, 2008.

can be easily dealt with in our cases. The analysis that there is an inconsistency is in fact based on assuming two inconsistent answers to the problem of scope. The inconsistency disappears once we settle on one answer.

The verdict that you ought not invest in all of your friends' startups relies on applying the respective precautionary principle to the extended activity. However, the verdict that you ought to invest in each of your friends' startups relies on your *only* applying the precautionary principle to each incremental activity separately: In that case, the precautionary principle does not constrain you, and you are free to choose what an ordinary weighing of costs and benefits would recommend – and I grant in that case you ought to engage in the incremental activity. If you were to *also* apply the precautionary principle to the extended activity in addition to each incremental activity, this might very well constrain what you should do on each day, as the precautionary principle then places restrictions on what series of action you may perform. The perception of inconsistency thus arises from an inconsistent answer to the problem of scope. If we answer it consistently, the inconsistency in recommendation disappears. We can either apply it just to incremental activities, in which case it is not true that you ought not perform all incremental activities in a series. Or we could apply it also (or only) to the extended activity, in which case it is not true that you should perform each incremental activity. The question we need to answer is at what scope or scopes the principle should be applied.

To those who have recognised the problem that risks which we are genuinely concerned about are often the result of taking many small risks that are individually judged too insignificant to trigger precaution, it has often seemed obvious that precautionary principles should be applied to extended activities. Jeryl Mumpower, for instance, writes:

A level of risk that is not of concern in any single instance may be viewed quite differently if it is part of an ongoing cumulative series. [...] This point underlines the importance of evaluating proposed *de minimis* (or any other risk management) schemes on the basis of the portfolio of risks that would accumulate over time from such a scheme, not on the basis of the apparent reasonableness of any single instance of its application.³⁴

Similarly, Katie Steele argues that we should “situat[e] any particular course of action within the broader field of actions that have similar consequences.”³⁵ Given the purpose of precautionary principles is to effectively guard against significant risks, and a risk is no less significant just because it is generated cumulatively, this is a natural response. For

³⁴Jeryl Mumpower. An analysis of the *de minimis* strategy for risk management. *Risk Analysis*, 6(4):437–446, 1986, p.442.

³⁵Katie Steele. The precautionary principle: A new approach to public decision-making? *Law, Probability and Risk*, 5(1):19–31, 2006, p. 22.

precautionary principles to do their intended jobs, it seems like they must be applied to extended activities.

While these authors suggest merely that precautionary principles should be applied at a large scope (only), in the deontology under uncertainty debate, Tenenbaum has suggested that deontological restrictions taking the form of precautionary principles, such as the requirement not to endanger the lives of others, should simply be applied to every scope of activity an agent is (considering) engaging in – including large ones.³⁶ After all, if I am a bad and frequent driver, driving badly on a particular day, and driving badly every day for a year are both things that I do. The precautionary principle could be thought of as potentially restricting all of my activities. If the latter endangers others, then the precautionary principle might demand that I don't drive at least on some days, even if no individual day of driving on its own would count as endangering others.

A common response to the problem of scope is thus that precautionary principles should be applied at a large, or alternatively even at all scopes. Indeed, precautionary principles could potentially be applied to all of an agent's risky lifetime choices taken together. Let me call this the 'global' scope. Precisely what the global scope for a policy-maker would be depends on questions of agential perspective that are often left vague in debates about precaution.³⁷ Are we specifying decision rules for regional, national or international policy-makers? And do we consider the policy-maker to have continuous agency throughout changes in government or constitution? Potentially, we could apply precautionary principles to all risky policy decisions facing a society for centuries to come all at once.

Suppose that the global scope was the or a right scope at which to apply precautionary principles. The problem now is that from that perspective, many kinds of precautionary principles that are standardly proposed in the literature are not plausible. Precautionary principles applicable to individual decision-making tend to state that even if an activity carries with it some very small, or non-negligible chance of a personal catastrophe (e.g. financial ruin, premature death or a severely harmful health outcome), then that activity should be avoided. Taking a global perspective, this principle would apply to the 'activity' of my entire life taken as a whole. But arguably any worthwhile way in which I could conduct my life will carry with it not only a very small or non-negligible, but a more substantial risk of personal catastrophe. Avoiding all non-negligible potential causes of premature death is not only practically difficult, if not impossible, it is also clearly undesirable. In the policy case, precautionary principles tend to state that even if an activity leads to a very small or non-negligible risk of catastrophe or serious harm to many people,

³⁶Sergio Tenenbaum. Action, deontology, and risk: Against the multiplicative model. *Ethics*, 127(3):674–707, 2017.

³⁷On this issue, also see Richard Bradley and Katie Steele. Making climate decisions. *Philosophy Compass*, 10:799–810, 2015.

then that activity should be avoided. But the normal course of civilisation over centuries viewed as a whole will bring with it not only a very small or non-negligible, but more substantial risks of various different kinds of serious harms. It is questionable whether a policy-maker should do everything in her power to make those risks negligible. For instance, adopting a science policy that would make the total risk of serious harms to human life and health from all scientific innovation expected to take place in the next 200 years negligible would likely be seriously restrictive to human progress.

The general point here is that viewed from the global scope, clearly more permissive likelihood thresholds are plausible than from smaller scopes. And the precautionary principles usually proposed appear to be calibrated so as to be plausible only from scopes smaller than the global scope. We may be able to formulate thresholds of harm and likelihood that are more plausible at the global scope, in terms of the lifetime risk of personal catastrophe that an agent should not cross. For instance, one such precautionary principle for individuals may state that, if I can and this does not come with exorbitant costs, I should avoid a lifestyle that generates more than a small risk of premature death. In the policy case, such ‘global’ precautionary principles would specify vague thresholds for the overall risk of various kinds of harm that a society should not cross viewed over the course of its existence.

While principles with more permissive likelihood thresholds would avoid the problem of implausibility at the global scope, they would look quite different from the ones usually proposed in the literature and policy discussions, which generally warn against much smaller risks of harmful outcomes. Moreover, they would only provide very limited action-guidance. The precautionary principle just proposed for individuals could be adhered to in a myriad of different ways. Even activities that unequivocally and substantially raise the risk of premature death may not be ruled out by such a principle. I could regularly ride a motorbike without a helmet, as long as I make efforts to substantially reduce all other risks of premature death. The kinds of precautionary principles that would be plausible from a global perspective would thus rule out some lifestyles, but probably wouldn’t rule out many particular kinds of activities.

Another way to avoid implausibility at the global scope would be to apply precautionary principles at large scopes short of the global scope, but not the global scope. Indeed, this might be what the authors suggesting that precautionary principles should be applied at ‘large’ scopes implicitly had in mind. But in that case where exactly should we stop? We can’t just instruct agents to pick the level of individuation that gives the intuitively ‘right’ result, as then the precautionary principle doesn’t do any action-guiding work. And it is hard to see what a principled general answer to that question could be. Without a principled answer, moreover, there is a danger that critics of precautionary principles have an easy time dismissing them by pointing to the implausible implications of the principles

when applying them to very large scopes.

To avoid these problems, I propose adding additional structure to how precautionary principles are typically formulated. There are at least two ways in which adding additional structure can help us respond to the problem of scope. The first is to add a clause that makes it the case that the agent is no longer constrained by the principle in those cases where the principle's recommendation would otherwise be implausible at the global scope. The most promising strategy, I think, would be to restrict precautionary principles to constrain agents only in cases where taking a risk of the magnitude specified in the principle is not required for achieving some much more important goal, such as, in the individual case, living a fulfilling life, or in the policy case, achieving technological transformation that significantly improves the lives of most people on Earth. Note that Tenenbaum's principle, which we labeled (c) above, included such a clause: "One may not endanger someone's life except in order to secure a very significant good."³⁸ If we include such a clause, the precautionary principle could be applied at all scopes after all, but would not restrict agents in cases where the upsides of taking the risk are especially important.

A second way in which additional structure can help address the problem of scope is to make the scope of activity the principles are meant to be applied to explicit in the formulation of the principle. The choice of time frame will to some extent be arbitrary, but should be chosen to be neither so short nor so long to lead to implausible recommendations. Consider, for instance:

PP-Risky Technology:* *If the use of a technology over the course of any ten-year period leads to a non-negligible risk of severe accident, and there are significantly safer and financially sustainable alternatives available, we should pursue one of these alternatives enough of the time to keep risks negligible.*

This principle implies, in our case, that we should not rely on the technology for any entire ten year period, but does not rule out using it for some of the time at least. This seems like the intuitively right recommendation for extended courses of action, and also avoids the problems we have previously discussed.³⁹

I think the first response is sensible if our sole purpose is to arrive at correct moral principles or principles of rationality, but not necessarily to find principles that are effectively action-guiding. The first response yields more general principles than the second; and it is independently plausible that precaution should not come at just *any* cost. On

³⁸Sergio Tenenbaum. Action, deontology, and risk: Against the multiplicative model. *Ethics*, 127(3):674–707, 2017, p.697.

³⁹The next section, however, will raise a challenge that, if unanswered, calls into question whether this principle can be rationally followed in transient cumulative likelihood cases like ours.

this view, precautionary principles allow us to disregard the potential upsides of a risky activity only within certain limits, thus ruling out what we might consider fanatical levels of precaution. When coming up with the right precautionary principles, we thus need to think systematically about what maximum cost precaution should come at.

However, I take the second response to the problem of scope proposed here to be superior for the purpose of effective action-guidance, for two reasons. First, on the first response, precautionary principles would only effectively guide choice if agents actually apply the principle at all scopes. Agents aiming to be guided by a precautionary principle in action would need to consider all of the activities – incremental, extended, or global – that they are, or are considering being engaged in, and check not only whether these involve risks that pass the harm and likelihood thresholds of a precautionary principle, but also whether taking such risks is necessary for some other important enough goal. There is thus a lot for agents to consider, a lot more room for disagreement, and for errors to be made if agents do not actively take a long-term perspective. Secondly, introducing another vague threshold in the formulation of a precautionary principle adds further to the more persistent problem of execution discussed in the next section.

The re-formulation in the second response considered here, by being explicit about scope, provides more straightforward and specific action-guidance. This is especially fitting in the policy arena, where precautionary principles are usually put forward for a specific policy or decision context, and are meant to be applied fairly directly to a particular policy problem at hand. Either way, however, the problem of scope has implications for how precautionary principles should be formulated. They must either include a clause that rules out implausibility at the global scope, or they should make scope explicit. Neither is commonly done in the literature on precautionary principles. And either way, precautionary principles end up being diachronic principles: In the second case explicitly so, as here precautionary principles explicitly constrain temporally extended courses of action. And in the first case, because agents are expected to apply the principle to all scopes of activity, including extended ones. In the face of risks that accumulate over time, precautionary principles can only hope to guard against excessive risk-taking when they are formulated and/or understood as diachronic principles. Next, I however want to raise two connected deeper challenges for such diachronic precautionary principles being effectively action-guiding in cumulative potential harm and cumulative likelihood cases.

V

Suppose an agent accepts the normative authority of a diachronic precautionary principle, which demands precautions regarding extended courses of action whose cumulative risk

passes some vague threshold. And suppose the recommended precaution needs to be implemented incrementally – as is the case when precaution takes the form of omission, reduction, or altered performance of some risky incremental activity. The agent now needs to make sure that the series of incremental activities she performs abides by the diachronic principle. In the kinds of cases we considered, which recommend reduction or omission, she needs to make sure not to perform a risky incremental activity too often.⁴⁰ But now the following reasoning suggests itself for any incremental activity:

Dangerous thought: Performing this one risky incremental activity will not keep me from abiding by the diachronic precautionary principle. Given this neutrality in terms of precaution, and its other benefits, I should thus perform it.

This thought is attractive for precisely the same reasons that we gave for the tolerance assumption above: Firstly, when thresholds are vague, small enough incremental activities cannot take agents from one side of a precautionary threshold to the other.⁴¹ For instance, investing in one more startup will never, on its own, make it the case that you violate a diachronic precautionary principle forbidding extended activities that come with a plausible chance of financial ruin. And second, in transient cumulative likelihood cases, past risks are continuously resolved, so new risky activities do not, at the time when they are performed, add to a growing likelihood of harm. For instance, where no accident has previously happened, any incremental decision to keep using the dangerous technology for another month will at that point be made against a ‘no risk’ background, and will thus, it seems, not take the agent close to the risk threshold for the extended activity.

The problem for the execution of diachronic precautionary principles is that if agents always reason along the lines of *dangerous thought*, for either reason, they will ultimately fail to abide by the diachronic precautionary principle. So unless we can show what is mistaken about always reasoning in these ways, or we can otherwise stop ourselves from engaging in this reasoning, diachronic precautionary principles cannot be rationally followed, calling into question not only their ability to effectively guide action, but also their validity.

⁴⁰Insofar as it arises from vagueness, the problem I point to here does not arise for precautionary principles where the recommended precaution is some separate measure that should be taken once risks resulting from our actions have passed the relevant thresholds, e.g. issuing warnings about drinking tap water once the water is polluted enough to pose health risks. In these cases, vague thresholds don’t create a temptation for indefinite delay, but only unclarity about precisely when the precaution should be taken. However, as proponents of precautionary principles tend to emphasise prevention and mitigation of risks over adaptation to and protection from potential harms, such less problematic cases can be expected to be rarer.

⁴¹Or, for the epistemicist, that they make such a difference is highly unlikely in the borderline region, and for the supervenientist, whether they do so is indeterminate in the borderline region (see Appendix). The same qualifications apply in the following.

Before addressing this challenge directly, let me note that there is some reason to think that versions of *dangerous thought* may be part of the explanation of why expressed commitment to precautionary principles has less successfully curtailed risks in some cases rather than others. Starting with the motivation for *dangerous thought* due to vagueness, Chrisoula Andreou has raised a similar concern for environmental policy contexts not explicitly featuring risk: Where incremental activities cumulatively bring about a vague harm, the thought that no incremental choice makes a difference may lead to policy procrastination or inability to prevent the harm. Global climate change in particular can be seen as a massive but creeping environmental problem. Andreou argues the thought of making no difference in incremental choices could be part of the explanation for the tragically insufficient action we have seen to date.⁴² Climate change is, of course, also a problem of risk management. Notably, the EU's commitment to the precautionary principle extends to tackling the risks of climate change.⁴³ Yet, as most countries in the world today, its current policies and pledges are insufficient for meeting the goals of the Paris climate agreement,⁴⁴ which has itself been criticised for being insufficiently precautionary.⁴⁵ The EU's stated commitment to the precautionary principle has had more tangible effects, for instance, when it comes to regulation of chemicals.⁴⁶ While a more careful analysis would of course be needed to establish this, the more tangible effects of such regulatory decisions may well be part of what explains the difference. Even if a policy-maker accepts the normative authority of a diachronic precautionary principle aiming to address the risks of catastrophic climate change, she might reason that no incremental policy decision and no day of delay ever makes a difference to whether she abides by the precautionary principle.

For a potential example of the *dangerous thought* at play in a case that features transient cumulative likelihood, consider Germany's decision, in 2010, to delay its phasing out of nuclear power by up to 14 years. The phase-out had previously been agreed on in the year 2000 under a different government, and would have involved closing the last nuclear power plants by 2021. The delay of the phase-out was itself quickly reversed in 2011 in response to public pressure after the Fukushima disaster. The last nuclear power plants in Germany are now set to close by 2022, close to the originally decided

⁴²Chrisoula Andreou. Environmental damage and the puzzle of the self-torturer. *Philosophy and Public Affairs*, 37(2):183–93, 2006 and Chrisoula Andreou. Environmental preservation and second-order procrastination. *Philosophy and Public Affairs*, 35(3):233–248, 2007.

⁴³European Commission (Environment Directorate General) and University of the West of England (Science Communication Unit). The precautionary principle: Decision-making under uncertainty. *Science for Environmental Policy: Future Briefs*, 18, 2017.

⁴⁴Climate Action Tracker. Country summary: EU <https://climateactiontracker.org/countries/eu/> (accessed August 6th, 2021).

⁴⁵Anju Sharma. Precaution and post-caution in the Paris Agreement: adaptation, loss and damage and finance. *Climate Policy*, 17(1):33–47, 2017

⁴⁶Mikael Karlsson. The precautionary principle in EU and US chemicals policy: a comparison of industrial chemicals legislation. In *Regulating Chemical Risks*, pages 239–265. Springer, 2010.

date.⁴⁷ Opposition to nuclear power has a long history in Germany, and (while this is a controversial application of precautionary reasoning) the precautionary principle was appealed to in justifying the decision to phase out, motivated, amongst other things, by the risk of accident.⁴⁸ Interestingly, the conservative-liberal coalition deciding on the extension of the phase-out did not disagree with the fundamental decision to phase out, but described nuclear power as a “bridge technology” in the more general transition to green energy.⁴⁹ What is interesting from our perspective is that nuclear power was already treated as such a bridge technology when the original decision of a gradual phase-out with a similar (roughly twenty year) time horizon was made in 2000, coupled with measures to boost green energy. With some, but insufficient progress made on the transition to green energy, the potential benefits of nuclear power as a bridge technology remained ten years later. And by 2010 we knew no major incident had happened at a German nuclear power station in the meantime. Those past risks were thus bygone (even if those from nuclear waste were not). This is where the *dangerous thought* motivated by the transiency of risks of accident potentially comes in: In these respects, looking forward, the calculation in 2010 looked very similar to that of 2000. Had it not been for the effects of the Fukushima disaster, the German nuclear phase-out would very likely have been delayed, in violation of the timeline that was considered to be called for by the precautionary principle in 2000.

So what can we say to an agent who accepts the normative authority of a precautionary principle, but is, or knows she will be, tempted by the *dangerous thought*? For one, such an agent could, in anticipation of this reasoning, enter a binding commitment to an extended course of action that abides by the precautionary principle, and thus remove the choices that may be affected by the *dangerous thought*. Or, if we think that even non-binding plans and resolutions convey rational pressure to go through with them, we may say that making a plan is enough to undermine *dangerous thought*: Yes, this incremental choice does not make me violate the precautionary principle, but I should not perform it because I planned not to.⁵⁰ In either case, abiding by the precautionary principle involves committing to a particular course of action that abides by the precautionary principle, and then going through with this plan. Given the vagueness of thresholds, this plan will to some extent be arbitrary.

Given how tempting the *dangerous thought* is, in practice it seems clear that plans

⁴⁷David Jacobs. The German Energiewende – history, targets, policies and challenges. *Renewable Energy Law and Policy Review*, 3(4):223–233, 2012

⁴⁸Ethikkommission Sichere Energieversorgung. Deutschlands Energiewende – Ein Gemeinschaftswerk für die Zukunft. Technical report, 2011.

⁴⁹Alexander Glaser. From Brokdorf to Fukushima: The long journey to nuclear phase-out. *Bulletin of the Atomic Scientists*, 68(6):10–21, 2012

⁵⁰These two responses correspond, roughly, to what two choice strategies from dynamic choice theory, sophisticated choice and resolute choice, would respectively recommend. See, for instance, Edward McClennen. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press, 1990 on these strategies.

and resolutions, or what Chrisoula Andreou in a similar context calls ‘implementation intentions’, as well as accountability mechanisms will be important for making sure that agents abide by diachronic precautionary principles. What this shows, at the very least, is that in the kinds of cases we have been discussing, precautionary principles are not effectively action-guiding on their own: Other decision-making tools are necessary. But I also think that pointing to binding or non-binding commitment devices is unsatisfactory unless we can say more about what is mistaken about always reasoning along the lines of *dangerous thought*. For one, in practice there are limits to the extent in which individuals and policy-makers can make truly binding commitments. More importantly, it would be odd indeed if, were it not for binding or non-binding pre-commitment devices, agents would be rationally required to always reason along the lines of *dangerous thought*, and thus irrational to abide by diachronic precautionary principles in our cumulative risk cases. This would speak of a deep tension between the principles and rational choice, and ultimately call into question the validity of the principles. We need to provide reasons to think that it is at least sometimes rationally permissible not to reason along the lines of *dangerous thought*, independently of a specific plan the agent has previously made.

I think a good case can be made for such a permission in cumulative potential harm and aggregating cumulative likelihood cases, when the *dangerous thought* is motivated by vagueness. We can see from the outset that *dangerous thought* will be attractive at any moment in time in cumulative risk cases, such as the investment case. But we can also see from the outset that you can only abide by the relevant diachronic precautionary principle if you refrain from performing the risky incremental activity enough of the time, that is, if you say ‘no’ to enough of your friends. If you accept the normative authority of the precautionary principle, it thus seems natural to also accept that there are permissions to refrain from incremental activities, even if *dangerous thought* makes them attractive. What grounds these permissions is your goal to keep risk below a certain vague level. While you know that no incremental choice can make you cross the threshold, there is still a sense in which each time you refrain from one, you contribute to this goal. Each time you say ‘no’ to a friend, you refrain from raising the cumulative potential harm. Likewise, each time you practice, you reduce the risk of failing your exam and thus contribute to remaining below the precautionary threshold. You can also, at any point when *dangerous thought* seems tempting step back and consider what potential series of choices would make you abide by the precautionary principle, and appreciate that some of the permissible courses of action involve you refraining from the particular incremental activity you are currently considering – though others may involve you performing this one and refraining from others. For these reasons, it seems plausible that it would at least be rationally permissible for agents who accept diachronic precautionary principles with vague thresholds to resist *dangerous thought* enough of the time.⁵¹ Commitment devices may then help them make

⁵¹This response is inspired by Diana Raffman and Sergio Tenenbaum’s solution to Warren Quinn’s

sure that they actually take enough of these permissions.

While I think this response is persuasive in the cumulative potential harm and aggregating cumulative likelihood cases, it is doubtful whether it applies to transient cumulative likelihood cases. The problem is that, in those cases incremental activities do not in the same way contribute to staying below the precautionary threshold, because the risks of harm stemming from each individual choice do not aggregate – rather, they are resolved one after the other. If bygone risk does not count, then each risky incremental choice really does only introduce an incremental risk against a no risk background. And this risk will itself be resolved before the next choice is made. From the perspective of the agent making an incremental choice and who disregards bygone risks, it is thus hard to see how incremental choices could be seen as making a contribution towards remaining below a certain cumulative risk level over time.⁵²

Our previous reasoning does extend, I think, if we continue to take into account risk even after it is bygone. If the risks we have already faced are something we continue to care about, then new independent risks we continue to take do add to an increasing stock of risk, as it were. We can then think of each incremental activity as adding to this stock, which may eventually pass a vague precautionary threshold. This could plausibly ground a permission to sometimes resist *dangerous thought*. I don't think it is irrational to care about bygone risk. Indeed, there is precedent in the economic literature of appealing to such a concern to show how those who violate expected utility theory can avoid time inconsistency.⁵³ And during the Covid-19 pandemic, citizens in many countries have been encouraged to think of activities that risk the spread of the virus as using up a “risk budget”, despite the transient nature of the risks imposed by the separate activities.⁵⁴ But I also think that most of the time, when people are concerned about risk, they are concerned about it in an exclusively forward-looking way. For a policy-maker tasked with acting on behalf of others in particular, there appears to be something problematically fetishistic about caring about bygone risks.

Self-Torturer Problem. Sergio Tenenbaum and Diana Raffman. Vague projects and the puzzle of the self-torturer. *Ethics*, 123(1):86–112, 2012. Warren Quinn. The puzzle of the self-torturer. *Philosophical Studies*, 59(1), 1990. They argue that vague goals that can only be achieved in series of acts create both diachronic requirements to perform some series of actions that achieves the goal, as well as permissions not to always act in accordance with the preferences that seem rational from a local perspective when considering to perform the next choice. Tenenbaum expands on this in his recent Sergio Tenenbaum. *Rational Powers in Action: Instrumental Rationality and Extended Agency*. Oxford University Press, 2020.

⁵²Johann Frick makes a similar point in a recent discussion of a paper by Joe Horton on PEA Soup. Johann Frick. Joe Horton's ‘aggregation, risk, and reductio’: Critical précis. *Ethics Discussion at PEA Soup*, August 2020.

⁵³See Mark Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, 1989.

⁵⁴See, for instance, Alyssa Bernanke. What is your risk budget? <https://covid-101.org/science/what-is-your-risk-budget/> (accessed August 6th, 2021).

So transient cumulative likelihood cases pose a real challenge. For agents who do not care about bygone risks, it is not clear what would ground a rational permission to resist *dangerous thought*. Those agents may still use binding or non-binding commitment mechanisms to force themselves to abide by diachronic precautionary principles. But in the absence of an argument in favour of a permission to resist *dangerous thought*, this would be a kind of commitment they know they would later rationally regret having made. And that would reveal the kind of conflict between rational choice and diachronic precautionary principles that calls into question the applicability of precautionary principles to these kinds of cases. I think this is a conclusion that should be troubling to proponents of precautionary principles.

VI

When risks are hard to quantify or precise quantification is hard to agree on, traditional forms of policy evaluation and rational choice cannot straightforwardly be applied. There is then a risk of either ignoring such risks or putting off taking precautionary action. The hope of proponents of precautionary principles is that they can be effective action-guiding principles that help agents address such risks. To play this role well, precautionary principles not only need to make plausible recommendations, they also need to provide useful action-guidance and effectively constrain choice.

This paper presented an important obstacle for precautionary principles serving their function well. Risks that are sufficiently likely and sufficiently harmful to trigger application of a precautionary principle often accumulate over time as the consequence of many incremental activities none of which creates significant risks when considered in isolation. Consequently, if we want precautionary principles to help us guard against those risks, they must be able to govern an agent's decision-making over time. To do so without making implausible recommendations, they must be understood to be diachronic principles which have some additional structure to how they are usually formulated: Either they must make explicit the scope of the activity they are meant to be applied to (which provides the more straightforward action-guidance), or they must introduce a further condition chosen to avoid implausibility when the principle is applied at a global scope.

But, as we saw in the last section, diachronic principles with vague thresholds are difficult to abide by for agents who make series of consecutive incremental choices. For every incremental choice, it is tempting to reason that this choice will make no difference to whether the diachronic principle will be violated, either due to the vagueness of thresholds, or because previously incurred risks are already bygone. At the very least, this limits the extent to which precautionary principles can be effectively action-guiding

in the face of cumulative risk. While diachronic precautionary principles might serve as plausible normative principles in cumulative risk cases, and ones that might garner more widespread support than any particular risk-cost-benefit analysis, they are unlikely, on their own, to be effective decision-making tools to guard against (policy) procrastination, inaction or recklessness. They need to be complemented with implementation intentions and commitment devices. If we want to bring about more precautionary policies, or foster greater precaution in our individual lives against temptations to the contrary, we are well advised not to place all our focus merely on formulating, promoting and accepting precautionary principles. More than precautionary principles is needed for a sense of urgency to be translated into action.

In cases where risks are transient, that is, when risks are continuously resolved before further incremental choices are taken, the worry in fact runs deeper, and it is not clear if precautionary principles have any purchase on such risks at all. This is a troubling implication, given that many of the risks policy-makers and individuals want to guard against are at least to some extent transient. The alternative would be to return to traditional risk-cost benefit analysis and expected utility theory, which are not to the same extent susceptible to tension between long-term and short-term perspectives.⁵⁵ But this would bring us back to the problems that motivated precautionary principles in the first place. We would then need to find other ways of overcoming these.

Appendix: Theories of Vagueness and the Application of Precautionary Principles to Incremental Activities

Can the major theories of vagueness accommodate the intuitively plausible tolerance assumptions in cumulative likelihood and cumulative potential harm cases? That is, can they accommodate the assumption that no sufficiently finely individuated incremental activity can make the agent in such cases cross the harm or likelihood thresholds specified in the antecedent of the relevant precautionary principle? And if they can't, do they imply that precautionary principles restrict incremental activities after all?

Contextualists about vagueness, such as Diana Raffman, can in fact accommodate

⁵⁵To some extent, this comes at the cost of an implausible implication: As Paul Samuelson first noted, the risk aversion we all exhibit for some small scale choices implies, within expected utility theory, implausibly extreme risk aversion at a larger scale (such as when considering compounds of many such small-scale choices). This is usually taken to imply that expected utility theory should be applied in a way that is virtually risk neutral at a small scale. Consistency between attitudes to small scale risky choices and larger compounds of them then comes at the cost of a failure to accommodate ordinary small-scale risk aversion. Paul Samuelson. Risk and uncertainty: a fallacy of large numbers. *Scientia*, 98:108–113, 1963. See also Matthew Rabin and Richard Thaler. Anomalies: risk aversion. *Journal of Economic Perspectives*, 15:219–232, 2001.

tolerance assumptions in the form that they matter for our analysis. According to contextualists, subtle context shifts will make it the case that, while there are sharp thresholds, these never lie ‘where we are looking’, that is, in an area we are actively considering.⁵⁶ And so one kind of tolerance assumption does hold true: No incremental activity *she is actively considering* ever makes an agent cross the vague harm and likelihood threshold of a precautionary principle. But this kind of tolerance assumption is all we need for the judgement that, when incremental activities are considered in isolation, the precautionary principle will never restrict an agent: When an agent is actively considering only whether to make the next investment or to practise the next hour, it is in fact true that this will not make the difference to whether there is a plausible chance of financial ruin or a substantial risk of failure, and so the antecedent of the precautionary principle turns out false.

Epistemicists about vagueness, such as Timothy Williamson, on the other hand, do accept that there is some incremental activity that takes an agent from, e.g. a non-substantial to a substantial chance of failure, even as the agent is actively considering that incremental activity. Epistemicists hold that thresholds can only be vague in the sense that it is impossible for an agent to know which incremental activity does so.⁵⁷ Borderline cases, on this account of vagueness, are simply cases of impenetrable uncertainty. Epistemicism thus introduces uncertainty about whether the antecedent of a precautionary principle is true for any incremental activity, and thus about whether a restriction applies to the agent. Indeed, when choices are individuated very finely, as they are in our examples, epistemicism implies that it would be subjectively very unlikely for any individual incremental activity to make an agent cross a vague threshold. According to epistemicism, agents can say with certainty that one of an extended series of incremental activities will make them cross a vague threshold, making the precautionary principle bite for certain when an extended enough series of choices is considered. But any individual choice is unlikely to make the antecedent of the principle true.

Whether precautionary principles constrain incremental activities, according to epistemicism, thus depends on how we think agents should act when there is only a small chance a deontological restriction applies. On an extremely risk averse construal, whenever there is any chance the antecedent of the precautionary principle is true, the agent should assume the restriction applies. On any more liberal construal, epistemicism is compatible with the intuitive analysis of the cases we started out with, at least if the incremental activities are individuated finely enough (and thus the chance any individual choice will take the agent over the threshold is small enough): The precautionary principle does not rule out incremental activities individually, while it may constrain series of the same choices when applied to extended courses of action. And even under the most risk averse construal, where the agent assumes the restriction applies whenever there is any

⁵⁶Diana Raffman. Vagueness without paradox. *Philosophical Review*, 103(1):41–74, 1994.

⁵⁷Timothy Williamson. *Vagueness*. Routledge, 1994.

chance of the antecedent being true, the recommendation of the precautionary principle in the borderline region is unclear, due to the widely accepted higher-order vagueness concerning when the borderline region of the application of a vague predicate begins. For the epistemicist, this amounts to the question of which is the first incremental activity for which there is some chance it is the sharp threshold. When there is higher-order vagueness, this question, for the epistemicist, is itself one that is subject to uncertainty. Given such higher orders of uncertainty, it is a complicated matter for an agent to work out whether the restriction of a precautionary principle applies to her or not when considering an incremental activity. The answer, in theory, will depend on a theory of choice we simply do not have yet, which tells us how to act when there is higher-order uncertainty about whether a deontological restriction applies to us. And in practice, on the epistemicist account of vagueness, the precautionary principle does not seem to have any useful action-guiding role to play when applied at the level of individual choices. Given the application of the precautionary principle now appears to involve making judgements under deep uncertainty, the precautionary principle at least no longer has a claim to being less subject to paralysis in the face of deep uncertainty than expected utility theory or risk-cost-benefit analysis in the original first-order choice context.

A third prominent theory of vagueness, namely supervaluationism, as defended by, e.g. Kit Fine, characterises borderline cases in such a way that for them, it is indeterminate whether the vague predicate applies or not.⁵⁸ The application of precautionary principles to incremental activities in the borderline region is complicated on such accounts, and our analysis will be parallel to the case of epistemicism. According to supervaluationism, there are many admissible precisifications of a vague predicate in the object language, and on each precisification, it is either true or false that the vague predicate applies. A sentence is supertrue if and only if it is true on all admissible precisifications. And its truth is indefinite if and only if it turns out true on some and false on others. Now take some incremental activity in the borderline region. There is likely one admissible precisification under which it is true that it takes you, e.g., from a chance of non-financial ruin to a chance of financial ruin. But there are going to be many others on which this is false. So the truth of the antecedent of the precautionary principle will be indeterminate. What does the precautionary principle recommend in such a case? If we understand the precautionary principle in the object language, it is going to be indeterminate if the agent is constrained by the principle. Under most precisifications, the agent is unconstrained by it, but under at least one, she is.⁵⁹ It is unclear how agents should act when it is indeterminate whether

⁵⁸Kit Fine. Vagueness, truth and logic. *Synthese*, 30(3-4):265–300, 1975.

⁵⁹We could, at this point, formulate the precautionary principle in the metalanguage. Two options here are: ‘If it is not superfalse that some choice has a plausible chance of leading to financial ruin, don’t do it’; and ‘if it’s supertrue that some choice has a plausible chance of leading to financial ruin, don’t do it.’ But these principles are themselves only going to give definite recommendations when there is no higher-order vagueness. And it is generally accepted that the borders of borderline regions of vague predicates are themselves vague.

some deontological restriction applies, and I am not aware of any account of choice under this kind of indeterminacy. As with epistemicism, we at best get the result that the precautionary principle does not give useful action-guidance when applied to incremental activities. And, unless we take a very conservative approach to choice under indeterminacy, this account, too, is compatible with what I take to be the intuitive analysis of these cases: The precautionary principle does not constrain agents when applied individually to sufficiently small incremental activities, while it does constrain a sufficiently long series of those activities when applied to the series.