Studies in Ethics, Law, and Technology

Volume 3, *Issue* 2 2009 *Article* 2

Risk, Precaution, and Emerging Technologies

Fritz Allhoff*

*Western Michigan University, fritz.allhoff@wmich.edu

Risk, Precaution, and Emerging Technologies*

Fritz Allhoff

Abstract

This paper explores a framework for thinking about risks inherent in emerging technologies; given uncertainty about the magnitude—or even nature—of those risks, deliberation about those technologies is challenged. §1 develops a conceptual framework for risk, and §2 integrates that conception into cost-benefit analysis. Given uncertainty, we are often pushed toward precautionary approaches, and such approaches are explored in §3. These first three sections are largely literature review, and then a positive argument for how to think about the relationship between risk, precaution, and uncertainty is offered in §4.

KEYWORDS: risk, precautionary principle, emerging technologies

*Department of Philosophy, Western Michigan University. Part of the writing of this paper was undertaken while a research fellow at The Australian National University's Centre for Applied Philosophy and Public Ethics; I thank ANU and CAPPE for their support. An earlier version of this paper was presented at a conference hosted by the University of Delaware—Environmental Nanoparticles: Science, Ethics, and Policy (2008)—and I thank those participants for useful discussion. Finally, I appreciate comments from Mark Cutter and Patrick Lin of this journal, as well as an anonymous reviewer who provided extraordinary feedback.

New technologies are constantly emerging, as well as new applications for those technologies. Alongside these developments, commentators enumerate various associated risks; such risks could be specific (e.g., environmental, economic) or else more general (e.g., social, ethical).[1] But comparatively little conceptual work has been done on the very nature of 'risk': what does it mean for something to be a risk (or to carry a risk)? And how does the nature of risk integrate, most fundamentally, with rational deliberation? On this latter question, proposals are often made regarding cost-benefit analysis or precautionary principles, but there are various issues with these proposals.

First, regarding cost-benefit analysis, it is unclear how this framework is meant to deal with much of the uncertainty inherent with risk, whether uncertainty about the probabilities of that risk being realized or else uncertainty about what the risks actually are. And, second, regarding precautionary approaches, the theoretical commitments of such approaches are rarely made transparent. It is easy enough to find instantiations of precautionary principles, but more work needs to be done to understand what these have in common with each other and what underlying structural features all precautionary approaches share.

Finally, the relationship between cost-benefit analysis and precautionary approaches is one that needs elucidation; I think that false claims have been made about these being "alternatives" to each other. This paper aims to clarify the aforementioned issues, both by undertaking a substantial review of the literature (§§1-3) as well as by offering new arguments (§4). This theoretical discussion floats free of any particular technology and instead offers a general platform by which to understand risk and precaution. This generality, though, hardly makes the present project of less interest to those of us thinking about individual technologies; rather the level of generality on offer is precisely that which we can apply to more specific contexts.

1. Risk

In technical discussions, there are various different senses in which ‘risk’ is used; here, I will follow Sven Ove Hansson and discuss four.[2] First, risk can be some unwanted event which may or may not occur. So we could say that

---

[1] My own research in this regard has predominantly been in the social and ethical issues of nanotechnology, though this essay generalizes beyond nanotechnology. For readers interested in these more specific discussions, see Fritz Allhoff *et al.*, *Nanoethics: The Ethical and Social Dimension of Nanotechnology* (Hoboken, NJ: John Wiley & Sons, 2007). See also Fritz Allhoff and Patrick Lin (eds.), *Nanotechnology & Society: Current and Emerging Ethical Issues* (Dordrect: Springer, 2008). See also Fritz Allhoff, Patrick Lin, and Daniel Moore, *What Is Nanotechnology and Why Does It Matter?: From Science to Ethics* (Oxford: Blackwell Publishing, in press).

[2] Sven Ove Hansson, “Philosophical Perspectives on Risk”, *Techné* 8.1 (2004): 10.

environmental impacts are one of the risks of some emerging technology; this is to say that the technology may or may not have these impacts and, furthermore, that the impacts would be negative. Note that both of these features are important for attribution of risk. If the technology definitely had some specific impact, then we would more appropriately call it a consequence of that technology rather than a risk: the uncertainty is one of the features of risk. And, second, the impact needs to be a bad thing for it to be a risk; otherwise we would call it something else, such as a (potential) benefit. These points might be obvious, but should help us in trying to conceptualize risk. Second, risk can be the *cause* of an event which may or may not occur. If we say that some technology carries environmental risks, what we mean is that it either might cause, tends to cause, or will cause negative environmental impacts. This postulation of a causal mechanism is more committed than the first conception of risk.

The third and fourth conceptions of risk are quantitative, as opposed to qualitative. The third conception holds that risk is the probability of an unwanted event which may or may not occur. So imagine that someone asks about the risk is of some technology having a certain environmental impact. An appropriate answer here might be, for example, 10%. The first sense of risk treated the environmental impact itself as the risk, whereas the second treated the technology as the risk (i.e., that which caused the impact). This third conception, though, tells us how likely it is that some impact will be realized. Fourth, and similarly, we could talk about the *expected outcome* of unwanted events. So imagine that there are 100 fish in some river that we are going to purify using nanoparticles. Further imagine that those nanoparticles are toxic to the fish population, and that some of the fish will die through the purification. We do not know which fish will die, but, given various epidemiological studies, we might reasonably issue a projection of 20%. The risk, then, is 20 fish, in the sense that we expect to lose that many fish. On the third sense, we are given the likelihood that something will happen (e.g., as a percentage), whereas the fourth sense gives us an expected outcome (e.g., in terms of some number of units lost). This fourth conception is the standard use of ‘risk’ in professional risk analysis. In particular, “’risk’ often denotes a numerical representation of severity, that is obtained by multiplying the probability of an unwanted even with a measure of its disvalue…”[3]

Henceforth, it is this fourth conception that I shall be most interested in, though some of the other conceptions will also recur. There are various reasons to focus on this fourth conception; as already mentioned, it is the standard use in risk analysis. One advantage that it has is that it allows us to assess risks quantitatively, which helps make them commensurable with benefits. For example, if we can say that some remediation will lead to an expected loss of 20

[3] *Ibid*.

DOI: 10.2202/1941-6008.1078

fish, this loss can then be compared, somehow, to the benefits of the remediation, such as more long-term benefits for fish, cleaner water for a local township, and so on. On the first conception, the risk would just be "the loss of fish". The second conception acknowledges that the remediation will *cause* the loss of fish, and the third conception tells us the likelihood. The fourth conception, though, ties all of these things together, telling us what we can *expect* to happen, given the remediation. And this is why it is the most useful for risk analysis, even if we can speak of risk in the other three senses.

Now that we have various conceptions of risk, including our preferred one, we can try to think more about how decision-making relates to risk. Of course, one of the hallmarks of risk is that we do not know for sure what will happen, given some course of action. It is this lack of epistemic certainty that makes decision-making under risk philosophically and practically interesting. If we knew that some course of action had a set of determinate consequences, some of which were good and some of which were bad, then decision-making would be a lot easier. To be sure, we might disagree about how to weight those good and bad consequences, such as if some technology had a positive economic upshot while having a negative environmental impact. Some might think that the environmental consequences were worth it, while others might not; this problem will occur in any society with pluralistic values. In such a situation, we have to think about how to render positive and negative consequences commensurable, and we further need to establish some democratic (or other) process for adjudicating disagreement. But in the case of epistemic uncertainty, this problem is exacerbated by the epistemological one, which is to say that we not only have to deal with a plurality of values, but we also do not even have epistemic certainty what the consequences will be. The values problem, then, is common to either scenario and is therefore not endemic to our discussion on risk.

Logically, there are four epistemic situations that we can be in with regards to risk.[4] The first of these is that we know the probability of some negative outcome. Imagine, for example, a case of Russian roulette in which a bullet is placed in one of six chambers of a revolver. Here we know the probability of a bullet being discharged, which is 1/6. Call this decision making under known probabilities: someone makes a decision whether or not to fire the gun, knowing what the probability is that a bullet will fire. Contrast decision making under known probability with decision making under uncertainty, wherein we know the probability only with insufficient precision. Return to the gun scenario and imagine that, last week, I put either two or three bullets in the chambers, but I have forgotten how many. If I choose to fire this gun, then I am doing so without known probabilities for the risks.

[4] Sven Ove Hansson, "What Is Philosophy of Risk?", *Theoria* 62 (1996): 170.

Finally, think of an extreme case of decision making under uncertainty: decision making under ignorance.[5] The ignorance, though, could be of two different sources, either of which would compromise our ability to determine some expected outcome. First, we might have little to no information about some specific outcome. Again, return to the gun example. Imagine that I pick up someone else's gun, not having any information about how many bullets are in the chamber, and then contemplate firing it. Assuming that I do not look in the chamber (and cannot otherwise tell anything by weight), I then have no information about the probability of a bullet discharging: that probability could be anywhere from zero if all chambers are empty to one if all chambers are loaded. Second, though, we might not even know what the outcomes *are*, much less how certain they are. Consider asbestos, for example, which became increasingly popular in the late 19th century as insulation. Despite the fact that even the Ancient Greeks observed lung damage in the slaves who wove it into cloth, the proclivity of asbestos to cause lung damage was not widely noted until the 1920s.[6] When asbestos became prevalent, the adverse health effects were largely unknown altogether, not just the probabilities that those effects would occur. It was not just that certain specific effects (e.g., mesothelioma and asbestosis) were unknown, but it was not even common knowledge that anything bad would happen to those who inhaled asbestos.

Technically, these three situations are all instances of decision making under uncertainty, though it is useful to think about the different variants in that regard. Putting all four together, then, here are the epistemic situations we can have in relation to risk:

1. Decision making with full knowledge of outcomes and probabilities;
2. Decision making with full knowledge of outcomes and some, though not all, knowledge of probabilities;
3. Decision making with full knowledge of outcomes and no knowledge of probabilities; and
4. Decision making with incomplete knowledge of outcomes (as well as their associative probabilities).

[5] See Sven Ove Hansson, "Decision Making under Great Uncertainty", *Philosophy of the Social Sciences* 26.3 (1996): 369-386.

[6] See, for example, W.E. Cooke, "Fibrosis of the Lungs Due to the Inhalation of Asbestos Dust", *British Medical Journal* 2 (1924): 147-150. In 1899 a London doctor, H. Montague Murray, connected the death of a factory worker to asbestos inhalation, after doing a post-mortem examination. The Cooke paper, though, as well as a report that came out shortly thereafter, were what established widespread recognition of the link. For the report, see E.R.A. Merewether and C.W. Price, *Report on Effects of Asbestos Dust on the Lung*, H.M. Stationery Office (1930).

DOI: 10.2202/1941-6008.1078

Again, (2)-(4) are all instances of decision making under uncertainty and (3)-(4) are both instances of decision making under ignorance; what separates these from each other is not their formal relationship, but rather the degree (or type) of uncertainty.[7]

When dealing with emerging technologies, in which epistemic situation are we likely to find ourselves? A ready observation is that it almost certainly is not the first one. The only time that we have full knowledge of outcomes and probabilities is likely to be in sorts of idealized cases, such as when we are talking about rolling dice or flipping coins that are known to be fair.[8] In reality, epistemic uncertainty is sure to abound. When trying to decide whether to pursue some course of action—especially more complex ones, like policy decisions—there will almost certainly be some negative consequences that may or may not follow, and it is very unlikely that we will have epistemic certainty either what those relevant probabilities would be or else even what the relevant consequences are. Of course, we at least hope to know the latter, and we also hope to know the former within some reasonable range of error. Whether this is true with risks in some particular emerging technology remains to be seen, though there is little reason to be optimistic given the often unknown risks of those technologies, or at least a wide range of uncertainty regarding the probabilities of those risks.

So what do we do with the uncertainty that we almost certainly face? Hansson, following Charles Sanders Peirce, offers an account of "uncertainty reduction" (cf., Peirce's "fixation of belief").[9] Hansson proposes that we reduce decision making under unknown probabilities to decision making under known probabilities, or even to decision making under certainty. For example, imagine that we are trying to figure out whether it will rain tomorrow. Various meteorologists get together and they all come up with estimates as to the likelihood of rain. Just for simplicity, suppose that three camps converge on reasonably close estimates: 70%, 80%, and 90% chance of rain. We must now make a decision about our day that hangs on whether it will rain (e.g., whether to plan a picnic). Further suppose that this is an instance of (2) above; we know what the outcomes are, but we do not have epistemic certainty as to the probabilities since the meteorologists disagree. What we will probably do is look at the testimony and aggregate it in some manner, thus, psychologically, abrogating the uncertainty. So, for example, we might take the testimony on board and then take the likelihood of rain to be 80%, effecting something like an

---

[7] Technically, everything could be classified as decision under uncertainty, so long as zero and one were allowed as the probabilities that some consequence would attain.

[8] *Ibid*., 171.

[9] *Ibid*., 172. See also Charles Sanders Peirce, "The Fixation of Belief" in Charles Hartshorne and Paul Weiss (eds.), *Collected Papers of Charles Peirce* (Cambridge, MA: Harvard University Press, 1934), pp. 223-247.

average of the reports. There are other things we could do, such as taking the median, picking our favorite meteorologist, excluding our least favorite meteorologist, and so on. But, pragmatically, we are certainly going to look to a way to reduce the uncertainty.

This reduction, seemingly, improves our epistemic status from (2) to (1). Of course, our actual epistemic status has hardly changed at all: we have not gained any more information, but have rather just adopted some strategy to convince ourselves that we know more than we do. Hansson thinks that we often take it a step further, moving ourselves toward known probabilities and then toward certainty. If we collapse the different testimonies to an 80% aggregation of rain, do we go on our picnic? Probably not: 80% is high enough that we convince ourselves that it *will* rain (i.e., that the chance of rain is 100%). And now our epistemic status is even better than (1) since we have full knowledge of the outcomes. Or so we would like to think; obviously we are still, actually, in (2).

What are we supposed to do with all this uncertainty? Some approaches, like Bayesianism, would have us assign probabilities to everything.[10] If we have no information at all, then maybe we just assign probabilities of 0.5; those prior probabilities will thereafter be revised as we start to garner evidence. In the long run, maybe these sorts of approaches will get it right, though they are not terribly practical, and they otherwise face short-term limitations. For example, imagine that some technological application may have some disastrous consequence, but we really have no idea whether it will. Should we proceed with the application? We could take the consequences, multiply them by 0.5, and then derive some expected cost; this expected cost can be compared to the expected benefits. But if, unbeknownst to us, the objective probability of the negative consequence is 0.9 (rather than our subjective 0.5), we could be really far off with our risk assessment. Of course, we could be off with it in the other direction, too, thus overestimating the risks rather than underestimating them. However, we might think that there is some sort of *asymmetry* between these sorts of errors: it is worse to be insufficiently cautious than it is to be overcautious. This sort of attitude gives rise to precautionary approaches, which will be presented in §3 and critically evaluated in §4. In the next section, though, let us take a step back and talk about cost-benefit analysis in general; the relationship between cost-benefit analysis and precautionary approaches will receive further discussion in subsequent sections.

---

[10] For an accessible introduction to Bayesianism, see Peter Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science* (Chicago: University of Chicago Press, 2003), ch. 14. For a more technical discussion, see John Earman, *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory* (Cambridge, MA: MIT Press, 1992).

DOI: 10.2202/1941-6008.1078

2. Cost-Benefit Analysis

§1 was meant to have two upshots. First, I wanted to try to conceptualize risk: various conceptions were considered and a proposal was issued to focus on the expected outcome conception. Second, I wanted to highlight the central role that uncertainty plays in risk, including the various guises under which it can appear. Now I propose to consider how cost-benefit analysis can be applied to decision-making under risk, with particular emphasis on how it looks under conditions of uncertainty.[11] This emphasis will be used to motivate precautionary approaches, though I will then return to the relationship those approaches bear to cost-benefit analysis.

Imagine that we are considering whether to perform some action, say φ. If we knew that φ had good consequences G and bad consequences B, then we could just think about whether the net effect was positive or negative (i.e., whether G-B>0). There are a lot of challenges here: the consequences need to be commensurable, they probably need to be (at least somewhat) quantifiable, people might disagree on how to weight them, and so on.[12] But we can imagine stripped-down examples that elide all of these interesting features. Imagine that we are running a business and are considering some marketing plan for the new genetic diagnostics device that our company has just developed; furthermore imagine that the marketing plan would cost $10,000 to execute and would increase our sales by $20,000. Finally, imagine that there are no other marketing plans under consideration and that, given your fiscal cycles, the decision has to be made immediately (i.e., before any other marketing plans could be developed). In this case, it seems straightforward that we should effect the plan since the benefits outweigh the costs, there are no other alternatives to consider, there are none of the messy complexities mentioned above, and so on.

---

[11] For our purposes, various nuances and conceptions of cost-benefit analysis are largely unimportant, though there is an important literature in this regard. One of the most ardent defenders of cost-benefit analysis is Richard Posner; see, for example, his *Catastrophe: Risk and Response* (New York: Oxford University Press, 2004). Cass Sunstein has written extensively on this topic; see, especially, his *The Cost-Benefit State* (Washington, DC: American Bar Association, 2002) and *Risk and Reason: Safety, Law, and the Environment* (Cambridge: Cambridge University Press, 2004). Frank Ackerman and Lisa Heinzerling critique cost-benefit analysis in *Priceless: On Knowing the Price of Everything and the Value of Nothing* (New York: New Press, 2003). Sunstein offers a review essay of contemporary scholarship, including Posner (2004) and Ackerman and Heinzerling (2003) in "Cost Benefit-Analysis and the Environment", *Ethics* 115 (2005): 351-385. See also Kristen Shrader-Frechette, *Taking Action, Saving Lives: Our Duties to Protect Environmental and Public Health* (New York: Oxford University Press, 2007).

[12] See, for example, W. Kip Viscusi, *Fatal Tradeoffs* (New York: Oxford University press, 1993). See also Ackerman and Heinzerling (2003).

Now imagine that, unlike the epistemic certainties of that case, there is the sort of epistemic uncertainty postulated in (1) above: we have known probabilities, but not certainties. The marketing plan still costs $10,000 to execute, but there is a 40% chance that it will fail, thus eliciting no increased sales. There is a 60% chance that it will succeed, thus eliciting the $20,000 in increased sales. All other details are the same. What do we do now? We already know the costs with certainty (viz., $10,000), but there is uncertainty about the benefits. We therefore calculate the expected benefits, which are:

$$0.4 * \$0 + 0.6 * \$20{,}000 = \$12000$$

The $12,000 in expected benefits is greater than the $10,000 in actual costs so we are still justified in pursuing the marketing plan, even given the possibility of its complete failure. Cost-benefit analysis, then, works not only when we have certainty regarding outcomes, but also when we have uncertainty but known probabilities.

Again, there are numerous other complexities to the cost-benefit approach; some were mentioned above. Returning to our earlier example of the river purification project, imagine that 100 fish will be killed, but the local township will have cleaner drinking water. These sorts of assessments have myriad complexities. Some of them are empirical: how *much* cleaner would the drinking water be? Would this *matter* in any significant way, such as health outcomes? Again, how *much*? And then come the issues of commensurability and values: imagine that the purification, while killing 100 fish, will lead to a 10% decrease in the local incidence of a certain water-borne disease, giardiasis, while having no other demonstrable effects. Is this worth it? There are not general answers to these sorts of questions, though we will return to some of them below; I just want to acknowledge some of these complexities.[13]

But, for present purposes, let us press on with our discussion of uncertainty. As shown above in the marketing plan cases, cost-benefit analysis seems promising when dealing with either known outcomes or else with known probabilities. Known outcomes, though, are not instances of risk at all, and so are not germane to that discussion. Cases of known probabilities, as mentioned in §1, are only likely to occur in idealized cases, such as ones involving fair dice and coins. While known outcomes or probabilities constitute positive epistemic statuses, these are not the epistemic statuses in which we are likely to find ourselves. Rather, we are more likely to find ourselves in (2)-(4) above: uncertain probabilistic knowledge, no probabilistic knowledge, and/or incomplete knowledge about outcomes. What guidance can cost-benefit analysis offer us

[13] Sunstein (2002), pp. 153-190 offers more discussion in a chapter called "The Arithmetic of Arsenic".

DOI: 10.2202/1941-6008.1078

now? Return to the marketing example and make the parameters as follows: the plan still costs $10,000 to execute, and it will either increase our sales or it will not (known outcomes). Imagine there to be a 40-80% chance that the plan will succeed in increasing revenues and a 20-60% chance that it will not (i.e., such that the chances of success and failure sum to 100%); our experts just cannot agree on the proper assessments. As before, sales go up by $20,000 if the plan is successful. Do we implement it? It is hard to figure out what to say. We could try to effect the sort of uncertainty reduction discussed above: maybe we act as if the probabilities are in the middle of the ranges, thus there being a 30% chance that the plan will fail and a 60% chance that it will succeed. The expected outcome, then, is $12,000, which means that we should execute the plan. But there is something overly simplistic about this approach. For example, even though there was a 20-60% chance of failure, it hardly follows that the *actual* chance of failure is 30%; all we really know is that the probability falls somewhere within that range. The same is true with the probability of success. Maybe the actual chance of failure is 60% and the actual chance of success is 40%. In that case, the expected increase to sales is $8,000, which is less than the cost of the marketing plan, so it should not be pursued. So, unlike when we know the probabilities and such privileged epistemic status leads to infallibility, we could make the *wrong decision* by applying cost-benefit analysis (in the above way, at least) when the probabilities are uncertain.

It is even worse when we move from limited knowledge of probabilities to no knowledge of probabilities. Imagine that we are considering the marketing plan, but we just have *no information* whether it will succeed or fail; maybe the CEO calls in looking for an immediate decision while all the relevant advisors are indisposed. Should we pursue the plan? As mentioned above, we could just give arbitrary assignments of probability to each outcome, 0.5 being the most plausible in cases of full ignorance. So there is a 50% chance that it will succeed and a 50% chance that it will fail, with an expected outcome of $10,000. Since this is how much the plan cost, we are neither any better nor any worse off by pursuing it or not. But this is almost certainly the (objectively) wrong answer since *any* other probability assignment would give a deterministic answer about our course of action. It is worse yet again if we do not even know what the outcomes are. Imagine that, unbeknownst to us, the marketing plan infringes on copyrights held by another company, thus exposing us to legal liability. If there is a 60% chance of success and a 40% chance of failure, then we might put our expected outcome on $12,000, thus meaning that we should pursue the plan. But this would actually be a disaster because, once we release it, we get sued for $50,000. Obviously, if we do not have all the information regarding outcomes, our abilities to make good decisions can be compromised.

These epistemic situations—(2)-(4) from above—are the ones in which we are most likely to find ourselves, and we see how cost-benefit analysis can get the wrong answer in these cases. This is not to say that we should not use cost-benefit analysis; indeed, as we will see in §4, it is not obvious that there is even an alternative. Rather, the point is just to show how uncertainty challenges cost-benefit analysis. This should not be surprising as uncertainty challenges *any* decision making approach but, in these contexts, we might think more about how and when to move forward on decision-making. Return to the above example where the CEO needs a decision on whether to effect the marketing plan, and we do not have any information on its prospects. One thing we might consider is to delay the decision until we have more information. Or, in the case where there is some unknown lawsuit waiting in the wings should we pursue the marketing plan, maybe we should not pursue the plan until we have reasonably convinced ourselves that no lawsuits are likely, or that there are any other negative externalities. I shall critically evaluate these possibilities in §4, but now let us consider the sort of approach that they suggest: precaution.

3. Precautionary Principles

Cost-benefit analysis under uncertainty poses risks, namely the risk of making the wrong decision. If we could somehow reduce the uncertainty, then we would occupy an improved epistemic status and be correspondingly more likely to make the right decision. The most obvious way to get rid of uncertainty is to hold off on making the decision until we have better knowledge regarding probabilities and outcomes. For example, if there is uncertainty regarding the probability of some outcome, then we could do more research and try to reduce the uncertainty. If there are unknown outcomes, then we could take more time and try to make sure that we have uncovered all of them. Particularly when we risk substantial and negative consequences, we should be wary of making hasty decisions. To wit, we might adopt something like a "precautionary principle".[14] Part of the challenge with the precautionary principle approach is getting clear about exactly what such a principle says, and various formulations abound. Charitably, there definitely seems to be merit to a principle that says we should not act hastily given the potential for substantial and negative consequences. When we try to pin down the details, though, it gets somewhat more complicated. Precautionary principles have been offered in various contexts, with environmental applications

[14] Note that much of the literature refers to *the* precautionary principle, though I shall talk about *a* precautionary principle or else precautionary principle*s*. The reason is that there is hardly any sort of definitive statement of "the" precautionary principle, but rather many different formulations that bear various relations to each other.

DOI: 10.2202/1941-6008.1078

being the most common; the reasons for this emphasis will become more clear below.

For starters, let us look at some actual precautionary principles in the hopes that we can try to understand their key features. There are many different formulations, as codified in national laws or international treaties.[15] For present purposes, however, we can take the context of issuance, as well as details regarding the issuing bodies, to be largely irrelevant. Consider the following three examples, which are representative. The first comes from the 1992 Rio Declaration of the UN Conference on Environment and Development (Principle 15):[16]

> In order to protect the environment, the precautionary approach shall be widely applied by States according to their capabilities. Where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation.

Another formulation is the 1998 Wingspread Consensus Statement on the Precautionary Principle, which holds that "When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically."[17] And, finally, consider the European Union's Communication on the Precautionary Principle (2000):

> The Communication underlines that the precautionary principle forms part of a structured approach to the analysis of risk, as well as being relevant to risk management. It covers cases where scientific evidence is insufficient, inconclusive or uncertain and preliminary scientific evaluation indicates that there are reasonable grounds for concern that the potentially dangerous effects on the

[15] Sunstein (2005) argues that Europe has been more sympathetic to precautionary approaches whereas the US has defended cost-benefit analysis (p. 351). I am more interested in the philosophical underpinnings of the approaches than their applications, but this phenomenon bears notice. See also Sunstein (2002). See also Arie Trouwborst, *Evolution and Status of the Precautionary Principle in International Law* (London: Kluwer Law International, 2002). Finally, see Poul Harremoës *et al*. (eds.), *The Precautionary Principle in the 20th Century: Late Lessons from Early Warnings* (London: Earthscan, 2002).

[16] Available online at http://www.un.org/documents/ga/conf151/aconf15126-1annex1.htm (accessed June 23, 2008).

[17] Available online at http://www.sehn.org/wing.html (accessed June 23, 2008).

> environment, human, animal or plant health may be inconsistent with the high level of protection chosen by the EU.[18]

These formulations are so varied that it is not even immediately obvious what they all have in common. We are told of a precautionary approach, precautionary measures, and a precautionary principle, though it hardly seems clear what any of these entails; reading the full documents, rather than just these excerpts, is not much more help. Risk (or threat) resounds throughout the different formulations, as does lack of evidence or certainty. But how do these pieces fit together in any meaningful sort of way? And, furthermore, how can we use those pieces to yield a *generalized* precautionary principle that abstracts away from the particular language used in these cases? In other words, what is the *logical structure* of precautionary principles? Is there one that they all share?

Before moving forward, it is worth acknowledging that much of the discussion regarding precautionary principles takes place in environmental contexts. The reason, as should become clear, is that the environment is an especially complex system; this complexity then gives rise to a lot of uncertainty regarding risks. For example, consider the introduction of rabbits into Australia.[19] Rabbits first came to Australia in the late 1780s, though the population explosion is thought to date to 1859. That fall, a landowner living near Melbourne released twenty-four rabbits into the wild to simulate British hunts, and other landowners then followed suit. Within ten years, there were literally millions of rabbits in the wild, and as many as 600 million nationwide by the mid-1900s; there are myriad ecological reasons for this population explosion, including mild winters and widespread farming (i.e., availability of food). The effects on Australia's environment have been disastrous, from species loss to erosion. Various countermeasures have been employed, such as shooting, poisoning, and fencing. Most dramatic (and effective) was the intentional introduction of a myxomatosis, a disease fatal to rabbits; the disease caused the rabbit population to fall to approximately 100 million, though resistance eventually spread. Australia again has in excess of 300 million rabbits, despite the introduction of calicivirus—another biological measure—in 1996.

The upshot of this example is that apparently trivial and benign acts can have catastrophic consequences: the release of twenty-four rabbits led to a rabbit population of over 600 million with dire economic and environmental impacts. Furthermore, those consequences could be unpredicted (or unpredictable) given

[18] Available online at http://www.gdrc.org/u-gov/precaution-4.html (accessed June 23, 2008).

[19] For a further discussion of this account, see Tim Low, *Feral Future: The Untold Story of Australia's Exotic Invaders* (Chicago: University of Chicago Press, 2002). For a more general theoretical account of invasive species, see Julie Lockwood, Martha Hoopes, and Michael Marchetti, *Invasion Ecology* (Hoboken, NJ: Wiley-Blackwell, 2006).

DOI: 10.2202/1941-6008.1078

the best scientific and other theories available. In addition to the consequences being negative and substantial, they can also be *irreversible*.[20] Once those first rabbits were released, thus began an inexorable march to the present circumstances. This is not to say that things could not possibly have been any other way than exactly as they are today (i.e., the first round of hunters could have caught all of their prey, myxomatosis could have had a slightly different epidemiological trajectory, etc.), only that the impacts on the relevant environmental system have been so substantial that any sort of complete remediation of the problem is virtually impossible. And, pragmatically, aside from the hunters catching all/most of the first rabbits, some roughly similar cascade of events would probably already have been prefigured.

More generally, any intervention into a well-functioning and complex system can have profound (and often negative) consequences. By definition, complex systems have many parts that fit together in complicated ways. Affecting either the parts or the relationships among them can have implications for the other parts and their interactions. Furthermore, feedback cycles can multiply these effects. And, finally, complex systems are the most epistemically intractable. In such systems, we are almost certain to have limited knowledge about their proper functioning and, therefore, knowledge about how some intervention will affect that functionality. As in the case of the Australian rabbits, small perturbations can be disastrous. The environment is obviously one such system, but there are others. Perhaps most analogous is the human body, on which many emerging technologies bear (e.g., genetic technologies, nanomedicine,[21] and so on). Human enhancement continues to receive much attention and is another realm in which complex systems carry novel challenges.[22] The present point, though, is just to establish the particular risks that complex systems offer, given our limited knowledge about them. As intimated and exemplified above, many precautionary approaches derive in environmental contexts for exactly this reason. Now we have seen the motivation for precautionary principles: to recognize the potential for dramatic and irreversible

---

[20] The concept of irreversibility is hardly transparent, though we shall not pursue further discussion here. For some of the conceptual complications, see Cass R. Sunstein, "Two Conceptions of Irreversible Environmental Harm" (May 2008). *University of Chicago Law & Economics, Olin Working Paper No. 407*. Available at http://ssrn.com/abstract=1133164 (accessed June 6, 2009). See also Neil A. Manson, "The Concept of Irreversibility: Its Use in the Sustainable Development and Precautionary Principle Literatures", *Electronic Journal of Sustainable Development* 1.1 (2007): 1-15.

[21] See, for example, Fritz Allhoff, "The Coming Era of Nanomedicine", *The American Journal of Bioethics* (in press).

[22] See, for example, Patrick Lin and Fritz Allhoff, "Untangling the Debate: The Ethics of Human Enhancement", *Nanoethics: The Ethics of Technologies that Converge at the Nanoscale* 2.3 (2008): 251-264.

damage in complex systems and to appreciate the limited epistemic situations in which we are likely to find ourselves in regards to those systems. With this in mind, let us return to our discussion of the logical structure of precautionary approaches.

In doing this, let us consider work done by Neil Manson.[23] Manson develops an account of precautionary principles by first looking to see what sorts of generic features they share; he then considers what relationship those features have to each other. In doing so, he does not presuppose that there is a single and general precautionary approach, as those features might have different relationships (or even be different) in different formulations. Nevertheless, he thinks that there is at least something that the different formulations must have in common in order for them to be plausibly considered precautionary principles.

Manson argues that given some *activity*, which may have some *effect* on the environment, a precautionary principle must indicate some *remedy*.[24] I think this sounds right, though a couple comments are worth making. Note the use of 'may', which bears emphasis. Central to all precautionary approaches is the notion of uncertainty: if we knew what the consequences were, then we could just see whether the net effect was positive or negative. Even if we had known probabilities for the consequences, we could formulate an expected outcome, as we worked through in §2. But, if we have unknown probabilities or unknown outcomes, everything becomes more complicated; I used these unknowns to motivate the idea of precaution in the first place. So it is critical to precautionary approaches that there be unknowns, as is reflected by 'may' above; other weak modal language, like 'possible', would also be appropriate, though see further discussion below. Second, Manson frames his discussion explicitly in terms of the environment, but I think that it generalizes beyond that context; as mentioned above, there are other contexts in which we have the same salient features, and my discussion will apply to those contexts as well. In what follows, I will offer the discussion at this more general level.

In addition to this acknowledgment of activities, effects, and remedies, Manson then argues that all precautionary principles must share a three-part structure. The first part is the *damage condition*, which specifies some characteristics of the effect in virtue of which the precautionary approach is warranted. The second part is the *knowledge condition*, which specifies the state of scientific knowledge regarding the relationship between the activity and the

[23] Neil A. Manson, "Formulating the Precautionary Principle", *Environmental Ethics* 24 (2002): 263-272. See also Per Sandin, "Dimensions of the Precautionary Principle", *Human and Ecological Risk Assessment* 5.5 (1999): 889-907 and Carl F. Cranor, "Toward Understanding Aspects of the Precautionary Principle", *Journal of Medicine and Philosophy* 29.3 (2004): 259-279.

[24] Manson (2002), p. 265.

DOI: 10.2202/1941-6008.1078

effect. Finally, the third part specifies the *remedy*, which is the course of action that decision-makers should take vis-à-vis the activity. Putting this all together, all precautionary principles must share this structure: if the activity meets the damage condition and if the link between the activity and the effect meets the knowledge condition, then decision-makers ought to effect the remedy. This is a very general structure, leaving many possibilities for particular precautionary principles. For example, consider the damage conditions, which could characterize the relevant effects in any of the following ways, among others: serious, harmful, catastrophic, irreversible, destructive of something irreplaceable, reducing or eliminating biodiversity, violating the rights of future generations, and so on. Knowledge conditions could invoke parameters like: possible, suspected, indicated by precedent, reasonable to think, not certainly ruled out, not reasonably ruled out, etc. And remedies could be: bans, moratoria, postponements, research into alternatives, attempts to reduce uncertainty, attempts to mitigate the damage conditions, and so on.[25]

So, for example, we could say that if some effect is serious and possible given some activity, then we ought not to perform that activity. The damage conditions do not always scale in a simple way (i.e., in terms of increasing damage) but, to the extent that they do, as the damages become greater, then we might require improved epistemic status before avoiding the remedy. For example, imagine that the damages could be either "serious" or "catastrophic", the latter obviously being worse. And then imagine that knowledge conditions could be "possible" or "not certainly ruled out". If catastrophic harms are possible, then we might trigger the remedy more readily than we would were the harms to be merely serious. Since "not certainly ruled out" carries a higher epistemic threshold than "possible" (i.e., it requires us to have greater knowledge), we should apply that knowledge condition more readily to the catastrophic damages than to the serious ones, all else being equal.

But all else does not have to be equal: rather than adjusting the knowledge condition as the effects become more negative, we could also adjust the remedy. Keeping the knowledge condition the same, then, again think of whether the effects proffer serious or catastrophic harms. As the harms become more substantial, then the remedy can simply become more restrictive. For example, we could say that, if the harms are serious, then the activity should be postponed. Alternatively, we could say that, if the harms are catastrophic, then the activity should be banned. This is not to say that, in either case would the harms necessarily be realized—because of the uncertain relationship between the activities and the effects—just that, all else equal, the remedy should be sensitive to the damage condition. All this is to say (and somewhat contrary to Manson's

[25] These various possibilities are adapted from Manson (2002), p. 267.

presentation) that these three conditions are not completely interchangeable, but rather should be interrelated to each other such that the above comparative desiderata attain. Given two different harms, which should be adjusted, the knowledge condition or the remedy? It depends. In some cases, our epistemic situation might be fairly hard to improve, so we might then adjust the remedies as the harms look more severe. In others, it might be the case that the remedies are hard to move (e.g., for legislative reasons), so we might then adjust the knowledge condition.

There are other things worth discussing, though many of them take us too far afield. Let us nevertheless make a few more observations before moving on. First, note that the knowledge condition effectively amounts to a burden of proof issue between the would-be practitioners of the activity and its opponents.[26] As this condition becomes more stringent proponents of the activity have more work to do in terms of ruling out some negative effect of their activity. For example, and perhaps counter-intuitively, 'possible' is more stringent than 'likely' in the sense that it is easier to rule out some effect being likely than its being possible; we might be able to show that it is not likely that our nanoparticles will have some negative effect on the environment without being able to show that such an effect is impossible. Whether we are willing to proceed with some activity given a possible effect rather than a likely effect, as suggested above, probably has to do with what that effect is, as well as whatever other recourses are available to us vis-à-vis remedies. Second, the remedies postulated by precautionary principles are quite commonly bans on the corresponding activities, though this hardly need be the case; above, we saw a wide range of other available remedies. Bans might make particular sense as the effects become worse, but it bears emphasis that the precautionary approach is not committed in this way.

Having gone through much of this abstract and theoretical discussion, let us return to the examples of precautionary principles presented above in order to see how well this theoretical account holds up against actual principles.[27] Consider again Principle 15 from the 1992 Rio Declaration of the UN Conference on Environment and Development:

> In order to protect the environment, the precautionary approach shall be widely applied by States according to their capabilities. Where there are threats of serious or irreversible damage, lack of

---

[26] For more discussion, see Carl F. Cranor, "Asymmetric Information, the Precautionary Principle, and Burdens of Proof" in Carolyn Raffensperger and Joel Tickner (eds.), *Protecting Public Health and the Environment: Implementing the Precautionary Principle* (Washington, DC: Island Press, 1999), pp. 74-99.

[27] Though also see Cranor (1999) for more discussion.

DOI: 10.2202/1941-6008.1078

> full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation.[28]

The damage condition here is explicit, namely that the specified damages are ones that are "serious or irreversible". The knowledge condition is also apparent: lack of full scientific certainty. The remedy is the non-postponement of measures to prevent environmental degradation. Putting it all together, and simplifying some of the language: If the damages are serious or irreversible, and if we lack full scientific certainty that those damages will occur, then we should not postpone measures to prevent environmental degradation. These statements are hardly transparent, though recognition of the underlying logical structure is definitely useful. To make this even less abstract, return to our example about river purification with nanoparticles. If we cannot rule out (cf., lack of scientific certainty) the possibility of those nanoparticles destroying the biodiversity of the river (cf., serious and irreversible), then we should not postpone measures that would prevent those harms. Those measures could include, for example, preventing the use of the nanoparticles at all. Now that we have a well-formed conception of the precautionary principle, let us subject it to critical evaluation.

4. Evaluating the Precautionary Principle

In evaluating the precautionary principle, it will be useful to have a particular conception in mind. The account developed above gives us the logical structure of precautionary principles, and there is nothing inherently problematic with a formal proposal that, given some potential for damage and given some epistemic status regarding the causal links between some activity and that damage, we should then effect some remedy. Rather, it is when we start specifying the damage condition, knowledge condition, and remedy that substantive critiques are possible. The hazard of picking a specific precautionary principle, though, is that criticisms of it will not necessarily apply to other variants; those variants could have different features that immunize them from the criticisms. Aware of this hazard, we nevertheless propose to proceed by focusing on a particular conception, though we will offer discussion of alternatives as we proceed.

The specific principle that we will consider is that one that is most commonly discussed in the literature: the catastrophe principle.[29] This principle specifies the damage condition as catastrophic, as opposed to lesser damages, such as harmful or serious ones. Its knowledge condition specifies possibility,

[28] Available online at http://www.un.org/documents/ga/conf151/aconf15126-1annex1.htm (accessed June 23, 2008).

[29] See also Manson (2002), pp. 270-274. Note that Posner (2004) explicitly defends cost-benefit analysis even under prospective catastrophe.

which is comparatively permissive: a lot of effects are possible even if they are not, for example, likely. And, finally, the remedy is a ban. As mentioned above, bans are common remedies offered by precautionary principles even if they are not, strictly speaking, required by such principles. So let us specify the catastrophe principle as follows: if, given some activity, some catastrophic effect is possible, then we should ban the activity. This formulation is substantive enough to be evaluated (i.e., the constitutive parts are specified) while still being general enough that the following discussion cuts across various ways it could be further specified vis-à-vis the particular activities or effects. There are three broad sorts of criticisms that have been lodged against this formulation; we shall consider them in turn.[30]

The first criticism goes to the knowledge condition, particularly its (extremely) weak modal operator: possibility. On the catastrophe principle, mere possibility of some catastrophe is enough to produce a ban against some activity. It is possible, in at least some sense, that some emerging technologies could destroy the world.[31] Surely that is catastrophic; ergo, no emerging technologies. But what is the sense of 'possibility' that matters? It has to be something stronger than mere logical possibility: it is (logically) possible, for example, that our emerging technology could lead, tomorrow, to some catastrophe on some inhabited planet in the deepest recesses of some other galaxy. But surely this is not physically possible, if for no other reason than it could not get there fast enough. Rather, what we need is some sort of physical possibility or, even better, empirical possibility: things may be physically possible that are nevertheless not likely to happen (e.g., decreased entropy in some complex system). This sort of possibility at least forestalls straw man objections against the catastrophe principle.[32]

Still, though, empirical possibility is extremely weak: a *lot* of things are empirically possible. For example, consider the notion that the Large Hadron Collider (LHC) could destroy the world.[33] Is this likely? No. Does any reasonable scientific evidence suggest that it would happen? No. But is it (empirically) *possible*? Yes. So, on the catastrophe principle, we could not run the LHC. This seems like the wrong answer, though, particularly given the (extreme) unlikelihood that the negative consequences would be realized.

---

[30] See also John Weckert and James Moore, "The Precautionary Principle in Nanotechnology", *International Journal of Applied Philosophy* 2.2 (2006): 191-204.

[31] For broader discussion, see Nick Bostrom and Milan M. Ćirković (eds.), *Global Catastrophic Risks* (Oxford: Oxford University Press, 2008).

[32] For more discussion, see David B. Resnik, "Is the Precautionary Principle Unscientific", *Studies in the History and Philosophy of Biological and Biomedical Sciences* 34 (2003): 329-344.

[33] See, for example, Dennis Overbye, "Gauging a Collider's Odds of Creating a Black Hole", *New York Times* (April 15, 2008). Available online at http://www.nytimes.com/2008/04/15/science/15risk.html?em (accessed October 7, 2008).

DOI: 10.2202/1941-6008.1078

Proponents of such an approach, however, could point out that the *magnitude* of the catastrophe justified the triggering of the remedy (e.g., a ban) *despite* the low probability of the catastrophe. And there has to be at least something right in this sentiment. Consider, for example, two cases. In the first, something extremely bad is going to happen with a 1% probability and, in the second, something somewhat bad is going to happen with a 50% probability. Which scenario is better? It has to matter what the magnitudes of the bad effects are. Imagine that we could render them financially, just to make the conceptualization simple. The first case has a 1% chance of having US$1B in damage. The second has a 50% chance of having US$1M in damage. Even though the probability is lower in the second case, the expected damages are twenty times higher in the first case. Therefore, we cannot just look at the (low) probability and say that we should proceed regardless. But what if the probabilities are really low and the consequences really bad (cf., the LHC example)? From an expected outcome approach, it does not matter; these would just "cancel out", thus giving results commensurable with more moderate values.

This gives rise to a second worry about the precautionary principle, which is to identify its relationship to traditional cost-benefit analysis. I think that this relationship has been poorly understood, particularly insofar as the precautionary approach is sometimes characterized as an "alternative" to cost-benefit analysis.[34] To motivate this part of the dialectic, consider that the precautionary approach is either something new (vis-à-vis cost-benefit analysis) or else it is not. On the former, it is supposed to be problematic and, on latter, it is not even interesting. Starting with the latter, remember that the defender of the catastrophe principle owes us some account of 'possible', both in terms of what it means and why it matters. Following the above discussion, let us assume that it means something like "empirically possible" and it matters because, despite the low probabilities, the potential effects are catastrophic. This sounds perfectly plausible, but then it just says the same thing as cost-benefit analysis; cost-benefit analysis can certainly accommodate low probabilities of catastrophes in terms of formulating expected outcomes.

Another way to go is to say that the precautionary principle really is saying something different. For example, the defender of the precautionary approach might deny that the environment is or has some singular value, which is commensurable among other values. Given that there is some catastrophic risk to the environment—however unlikely—that risk just trumps all other considerations. This sort of line is different from cost-benefit analysis in the

[34] See, for example, Manson (2002), p. 264; Weckert and Moor (2006), p. 191. Sunstein (2005) alleges a "tension" between precautionary and cost-benefit approaches (p. 352) though then goes on to suggest that the views are "complementary" (p. 355). These certainly look like different claims, but I am ultimately sympathetic to the latter, as will be expressed below.

sense that the latter would allow us to consider the *benefits* of some activity, rather than merely having to stop at an identification of the risks. But, for a number of reasons, this has to be wrong. First, it allows extremely low probabilities to derail entire activities. (Again, one could point to the magnitude of the consequences, but then this just brings us back to the first horn of the dilemma.) Second, these low probabilities—which nevertheless establish *possibility*—could be effectively impossible to reduce to zero. Imagine we can show a 1% chance of some effect, or a 0.1% chance, or a 0.01% chance: in no case have we shown that it is impossible. If the precautionary approach is meant to do something different than cost-benefit analysis, then it would be paralyzing. Third, this is simply irrational. Imagine that, if we φ'd, there was an X% of some cost C; further imagine that C is really bad. Should we φ? It is impossible to even conceptualize this question without knowing benefits would attain by φ'ing (as well as their associative probabilities). Imagine there is some evil deity who asks for a tithe, lest he destroy the planet. Furthermore, imagine that he might destroy the planet anyway, given the (remote) probability that he finds the tithe unacceptable. So, if we tithe, then it is possible that he will destroy the planet. The defender of the catastrophe principle therefore has to say that we cannot tithe, even if the deity will certainly destroy our planet if we do not. This does not make any sense: it is completely irrational to allow remote risks to completely preclude our consideration of the associative benefits for some course of action.

The evil deity example gives rise to a third criticism of the catastrophe principle; this criticism holds not just that the principle is false, but rather that it is incoherent. Consider Cass Sunstein: "Because risks are on all sides of social situations, and because regulation itself increases risks of various sorts, the principle condemns the very steps that it seems to require."[35] So imagine that it is possible that some activity give rise to some catastrophe. Therefore, we ban that activity. But surely it is possible that *the ban* risks a catastrophe as well. So we cannot ban the activity. Return to our example about water purification using nanoparticles: this practice could (even if not likely) have disastrous effects on the environment. But a failure to have clean water could (and probably more likely would) lead to disastrous effects, particularly vis-à-vis the world's poor

[35] Sunstein (2005), p. 355; see also pp. 366-369. Sunstein means this criticism to apply to the precautionary principle more generally, rather than to the catastrophe formulation in particular. I disagree and think that the criticism, at best, attaches to catastrophe-like formulations because different knowledge conditions (e.g., ones requiring "likely" rather that "possible") are unaffected by the criticism. See also Cass R. Sunstein, "Beyond the Precautionary Principle", *Pennsylvania Law Review* 151 (2003): 1003-1058 and Cass R. Sunstein, *Laws of Fear: Beyond the Precautionary Principle* (Cambridge: Cambridge University Press, 2005). For a detailed response to the incoherence objections, see Jonathan Hughes, "How Not to Criticize the Precautionary Principle", *Journal of Medicine and Philosophy* 31 (2006): 447-464.

DOI: 10.2202/1941-6008.1078

who are increasingly without drinking water.[36] The effects on them directly are bad enough, but there could be added effects in terms of political destabilization, global conflict, and so on. The catastrophe principle would say that we cannot purify the water and, similarly, that we cannot effect the ban against the purification. In other words, it says we cannot φ and we cannot $\sim\varphi$. This is logically impossible, therefore the principle is incoherent. The incoherence charge is a strong one, and certainly one best avoided. For example, so long as one of the catastrophic effects is more likely than the other (e.g., as follows from φ or $\sim\varphi$), then maybe the advocate just guards against the most likely catastrophe. But this would require further emendation to the principle, and then risks some of the other criticisms presented above.[37]

Having seen various criticisms, let me now offer my own view.[38] I think that there are two fundamental issues with precautionary approaches. The first has to do with the knowledge condition. In the catastrophe formulation, mere possibility was enough to force the ban on some activity. Some people have wanted to say that this leads to bans too easily since negative effects will always be possible, even in our sense of empirical possibility. This does not worry me, though, because of the potential magnitude of those effects. If the probability of the effects is really low, but the negative consequences of the effects are really high, then we should take the risk seriously. Part of the problem is undoubtedly epistemic as we will not always know what the probabilities are, and we certainly cannot rule out that they are zero (as the catastrophe approach would seemingly require). I will return to that below, but suffice it to say, that unlikely but catastrophic risks should obviously play a part in our decision-making. We hardly need a precautionary approach, though, to tell us that; no reasonable person would deny it.

One obvious way to make the precautionary approach more permissive is to relax the knowledge condition. For example, we might say that the negative effects need to be, not just possible, but likely. This project becomes more epistemically tractable in the sense that it is easier to establish likelihood than it is to rule out possibility; this is not to say that establishing likelihood is easy, but ruling out possibility is extremely hard. Note that, pragmatically, this suggestion transfers the burden of proof from the proponent of some activity to its detractor. For example, we might not be able to rule out the possibility of the LHC

---

[36] See, for example, Allhoff, Lin, and Moore (in press), ch. 7.

[37] Another criticism, not presented here, is that precautionary approaches contribute to, and even promote, unfounded public fears. See Adam Burgess, *Cellular Phones, Public Fears, and a Culture of Precaution* (Cambridge: Cambridge University Press, 2004).

[38] For more detailed responses to some of these criticisms, see Stephen M. Gardiner, "A Core Precautionary Principle", *Journal of Political Philosophy* 14.1 (March 2006): 33-60. Gardiner defends a particular version of the precautionary principle, arguing that his formulation—different from the catastrophe principle—is immune to standard criticisms.

annihilating the world, but can it be proven to be likely? Defenders of precautionary principles think that the burden of proof should be on the would-be facilitator of some catastrophe; opponents claim that the principles are too restrictive. Where should the burden of proof go? I do not think that this question or corresponding conception is very useful. Rather, what matters are what the risks *are*. They might be hard to determine but, conceptually, the risks are what matter, not where the burdens of proof fall. From a procedural or regulatory perspective, burden of proof might be important, but there are ways of dealing with it (e.g., further research, independent commissions). But, philosophically, the focus should be on the risks themselves.

This, then, brings us back to the second fundamental issue with the precautionary approach, which is its relationship to cost-benefit analysis. As suggested above, I think that there has been a lot of confusion regarding this issue, particularly in claims that precautionary principles are alternatives to cost-benefit analysis. Cost-benefit analysis cannot possibly be wholesale wrong as an approach to decision making. In our everyday lives, we continually weigh costs and benefits (discounted by their perceived probabilities) and make decisions based on those assessments; such an approach is almost the paragon of rationality.[39] If you were facing a great but unlikely benefit versus a great and likely benefit for opposing courses of action, which would you pick? The answer is so trivial as to lead us to wonder what all this dialogue over the precautionary principle is supposed to contribute.

There seem to be two possibilities in this regard. The first is that there are certain domains in which the precautionary principle is supposed to supplant cost-benefit analysis. For example, consider environmental contexts in which serious and irreversible harm is possible; this is the sort of context in which we often see precautionary principles surface. But why would cost-benefit analysis be ill-equipped to handle this situation? Certainly cost-benefit analysis can accommodate concepts like 'serious' and 'irreversible' since these have obvious upshots in terms of risk assessment. It cannot be those concepts that activate the precautionary approach as an alternative to cost-benefit analysis. What about the environmental context itself? Maybe we should exercise extra caution when dealing with the environment because of the sort of thing that it is or the moral features that it has. Even if this is true, though, cost-benefit analysis would still work: only the *weighting* of the relevant considerations would change once we properly appreciated environmental values. In other words, imagine that we rally behind the precautionary approach because we really decide that the environment is important. What have we gained? We could just do cost-benefit analysis and

[39] A similar attitude is expressed by Posner (2004), who argues that cost-benefit analysis "is an indispensible step in rational decision making", even in under catastrophic risk (p. 139); quoted in Sunstein (2005), p. 363.

DOI: 10.2202/1941-6008.1078

maintain that environmental costs are really bad and environmental benefits are really good. If precautionary approaches effectively increase the weighing of environmental considerations, we could afford similar weightings through cost-benefit analysis.

Whoever wanted to defend the "supplanting" model would now have to argue that the environment is not simply one value among many—or even an important value, as cost-benefit analysis could surely accommodate—but rather that it is patently incommensurable with other values. So, the argument would go, we cannot use cost-benefit analysis because the environment is special and cannot be compared to other values. To figure out whether to destroy the Redwood Forest, we hardly focus on the joy that would be derived from the proposed theme park, or on how much money people would be willing to pay to access those trees as against the alternative theme park. This joy and the associative economic preferences are morally relevant, but these are incommensurable with the value of the forest.[40] Surely the forest must be preserved regardless. Or so the dialectic might proceed. However, it cannot be right that environmental values are incommensurable with others: imagine that terrorists will destroy the entire world unless we destroy a single tree. Save the tree? Forests and trees matter, but so do a lot of other things, and we have to have a complex value system that accommodates all of those values. For these reasons and those above, I therefore reject the idea that precautionary approaches are meaningful alternatives to cost-benefit analysis.

Rather, I think that precaution *supplements* cost-benefit analysis *given uncertainty*.[41] As we saw in §1, there are various epistemic situations in which we might find ourselves with regards to risk. If we know that some act A has an X% chance of realizing some benefit B while, at the same time, having a Y% chance of realizing some cost C, then we just compare X*B+Y*C with the alternatives to A and pick the best expected outcome. As I discussed in §2, this becomes more complicated when we do not know X or Y. It is even worse when we do not know B and C, either. Precaution is a risk-averse strategy for dealing with uncertainty.[42] If we know that there is an X%-Y% chance of some cost C,

---

[40] A classic on this issue is Mark Sagoff, "At the Shrine of Our Lady of Fatima; or, Why All Political Questions Are Not All Economic", *Arizona Law Review* 23.4 (1981): 1283-1298.

[41] Cf., Posner (2004). See Gardiner (2006) for a contrary proposal.

[42] Consider, for example, Heinzerling and Ackerman (2003), who mean to be offering a critique of cost-benefit analysis and a defense of precaution. If, for example, our nation spends more than it needs to on regulatory protection, its "preference is to tilt toward overinvestment in protecting ourselves and our descendents." (p. 227); quoted in Sunstein (2005), p. 359. But this "precaution" is just demonstrating a collective agreement that the prospective negative consequences are really bad and that we hardly want to countenance their actualization.

However, this is hardly antithetical to cost-benefit analysis, which has to be able to accommodate our preferences: what else would "cost" and "benefit" even *mean* as wholly precaution might, for example, tell us to act *as if* the probability were the higher value, Y%. And, if we were considering some uncertain benefit, we might act as if the probability were the lower value. But this then integrates quite well with cost-benefit analysis: it just requires us to be conservative in our assessments.

Whether we should be conservative does not depend on the (non-epistemic) values at stake nor their probabilities, which are treated straightforwardly through cost-benefit analysis. Rather, the conservativeness is dictated by the (epistemic) value of uncertainty and our predilections against it. The disvalue of uncertainty is hardly obvious; there are certainly contexts in which most of us prefer it (e.g., opening presents). When making decisions about moving forward with particular technologies, we just have to think about how tolerable uncertainty is, particularly given the potential consequences.

References

Ackerman, Frank and Lisa Heinzerling. 2003. *Priceless: On Knowing the Price of Everything and the Value of Nothing* (New York: New Press).

Allhoff, Fritz, *et al*. 2007. *Nanoethics: The Ethical and Social Dimension of Nanotechnology* (Hoboken, NJ: John Wiley & Sons)

Allhoff, Fritz and Patrick Lin (eds.). 2008. *Nanotechnology & Society: Current and Emerging Ethical Issues* (Dordrect: Springer).

Allhoff, Fritz, Patrick Lin, and Daniel Moore. In press. *What Is Nanotechnology and Why Does It Matter?: From Science to Ethics* (Oxford: Blackwell Publishing).

Bostrom, Nick and Milan M. Ćirković (eds.). (2008). *Global Catastrophic Risks* (Oxford: Oxford University Press).

---

independent of our preferences? There might be cases, like the forests, where we want to ascribe value independently of our preferences, but there are other cases, like whether to effect a tax increase to support a farm subsidy, that cannot be understood without thinking about how put off we would be by the tax increase, what impacts it would have on the food supply (and whether we would care about those), and so on.

So, when we "over"-invest, all we are doing is demonstrating that we take the cost to be worth the protection that it affords us against a negative outcome; this protection could be economic, psychological, moral, or symbolic. (Cf., for example, the war on terror, which almost certainly costs far more money than it could ever prevent in terms of economic damages.) This is not to say that we are infallible with all of our protective investments, though it is to say that we can rationally accommodate risk-aversion under a cost-benefit framework.

DOI: 10.2202/1941-6008.1078

Burgess, Adam. 2004. *Cellular Phones, Public Fears, and a Culture of Precaution* (Cambridge: Cambridge University Press).

Cooke, W.E. 1924. "Fibrosis of the Lungs Due to the Inhalation of Asbestos Dust." *British Medical Journal* 2: 147-150.

Cranor, Carl F. 1999. "Asymmetric Information, the Precautionary Principle, and Burdens of Proof." In Carolyn Raffensperger and Joel Tickner (eds.), *Protecting Public Health and the Environment: Implementing the Precautionary Principle* (Washington, DC: Island Press, 1999), 74-99.

Cranor, Carl F. 2004. "Toward Understanding Aspects of the Precautionary Principle." *Journal of Medicine and Philosophy* 29.3: 259-279.

Earman, John. 1992. *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory* (Cambridge, MA: MIT Press).

Gardiner, Stephen M. 2006. "A Core Precautionary Principle." *Journal of Political Philosophy* 14.1: 33-60.

Global Development Research Center. EU's Communication on Precautionary Principle. Available at http://www.gdrc.org/u-gov/precaution-4.html (accessed June 23, 2008).

Godfrey-Smith, Peter. 2003. *Theory and Reality: An Introduction to the Philosophy of Science* (Chicago: University of Chicago Press).

Hansson, Sven Ove. 1996. "Decision Making under Great Uncertainty." *Philosophy of the Social Sciences* 26.3: 369-386.

Hansson, Sven Ove. 1996. "What Is Philosophy of Risk?" *Theoria* 62: 169-186.

Hansson, Sven Ove. 2004. "Philosophical Perspectives on Risk." *Techné* 8.1: 10-35.

Harremoës, Poul, *et al.* (eds.). 2002. *The Precautionary Principle in the 20th Century: Late Lessons from Early Warnings* (London: Earthscan).

Hughes, Jonathan. 2006. "How Not to Criticize the Precautionary Principle." *Journal of Medicine and Philosophy* 31: 447-464.

Lin, Patrick and Fritz Allhoff. 2008. "Untangling the Debate: The Ethics of Human Enhancement." *Nanoethics: The Ethics of Technologies that Converge at the Nanoscale* 2.3: 251-264.

Lockwood, Julie, Martha Hoopes, and Michael Marchetti. 2006. *Invasion Ecology* (Hoboken, NJ: Wiley-Blackwell).

Low, Tim. 2002. *Feral Future: The Untold Story of Australia's Exotic Invaders* (Chicago: University of Chicago Press).

Manson, Neil A. 2002. "Formulating the Precautionary Principle." *Environmental Ethics* 24: 263-272.

Manson, Neil A. 2007. "The Concept of Irreversibility: Its Use in the Sustainable Development and Precautionary Principle Literatures." *Electronic Journal of Sustainable Development* 1.1: 1-15.

Merewether, E. R. A. and C. W. Price. 1930. *Report on Effects of Asbestos Dust on the Lung*, H.M. Stationery Office.

Overbye, Dennis. 2008. "Gauging a Collider's Odds of Creating a Black Hole." *New York Times*. Available online at http://www.nytimes.com/2008/04/15/science/15risk.html?em (accessed October 7, 2008).

Peirce, Charles Sanders. 1934. "The Fixation of Belief" in Charles Hartshorne and Paul Weiss (eds.), *Collected Papers of Charles Peirce* (Cambridge, MA: Harvard University Press), pp. 223-247.

Posner, Richard. 2004. *Catastrophe: Risk and Response* (New York: Oxford University Press).

Report of the United Nations Conference on Environment and Development. Available at http://www.un.org/documents/ga/conf151/aconf15126-1annex1.htm (accessed June 23, 2008).

Resnik, David B. 2003. "Is the Precautionary Principle Unscientific." *Studies in the History and Philosophy of Biological and Biomedical Sciences* 34: 329-344.

Sagoff, Mark. 1981. "At the Shrine of Our Lady of Fatima; or, Why All Political Questions Are Not All Economic." *Arizona Law Review* 23.4: 1283-1298.

DOI: 10.2202/1941-6008.1078

Sandin, 1999. "Dimensions of the Precautionary Principle." *Human and Ecological Risk Assessment* 5.5: 889-907

Science and Environmental Health Network. Precautionary Principle. Available at http://www.sehn.org/wing.html (accessed June 23, 2008).

Shrader-Frechette, Kristen. 2007. *Taking Action, Saving Lives: Our Duties to Protect Environmental and Public Health* (New York: Oxford University Press).

Sunstein, Cass. 2002. *The Cost-Benefit State* (Washington, DC: American Bar Association).

Sunstein, Cass. 2003. "Beyond the Precautionary Principle." *Pennsylvania Law Review* 151: 1003-1058

Sunstein, Cass. 2004. *Risk and Reason: Safety, Law, and the Environment* (Cambridge: Cambridge University Press).

Sunstein, Cass. 2005. "Cost Benefit- Analysis and the Environment." *Ethics* 115: 351-385.

Sunstein, Cass. 2005. *Laws of Fear: Beyond the Precautionary Principle* (Cambridge: Cambridge University Press).

Sunstein, Cass. 2008. "Two Conceptions of Irreversible Environmental Harm." *University of Chicago Law & Economics, Olin Working Paper No. 407* Available at SSRN: http://ssrn.com/abstract=1133164 (accessed October 2008)

Trouwborst, Arie. 2002. *Evolution and Status of the Precautionary Principle in International Law* (London: Kluwer Law International).

Viscusi, W. Kip. 1993. *Fatal Tradeoffs* (New York: Oxford University press).

Weckert, John and James Moore. 2006. "The Precautionary Principle in Nanotechnology." *International Journal of Applied Philosophy* 2.2: 191-204.