MEDICAL ETHICS

How to take deontological concerns seriously in risk–cost–benefit analysis: a re-interpretation of the precautionary principle

S D John

J Med Ethics 2007;33:221–224. doi: 10.1136/jme.2005.015677

In this paper the coherence of the precautionary principle as a guide to public health policy is considered. Two conditions that any account of the principle must meet are outlined, a condition of practicality and a condition of publicity. The principle is interpreted in terms of a tripartite division of the outcomes of action (good outcomes, normal bad outcomes and special bad outcomes). Such a division of outcomes can be justified on either “consequentialist” or “deontological” grounds. In the second half of the paper, it is argued that the precautionary principle is not necessarily opposed to risk–cost–benefit analysis, but, rather, should be interpreted as suggesting a lowering of our epistemic standards for assessing evidence that there is a link between some policy and “special bad” outcomes. This suggestion is defended against the claim that it mistakes the nature of statistical testing and against the charge that it is unscientific or antiscientific, and therefore irrational.

Firstly, a defence of the precautionary principle should provide clear criteria for when and how the principle is to be applied. Secondly, the principle should not be justified in terms of some contested conception of the good. Rather, in pluralist societies, an argument for the precautionary principle should constitute a Rawlsian “free-standing module”, on which there might be “overlapping consensus”.⁴ Most formulations of the precautionary principle are too vague to meet the first condition. Furthermore, at least in the original context of environmental policy, many formulations of the principle appealed to a contestable view of man’s relationship to nature, violating the second condition.⁵ My question, then, is whether, in public health contexts, we have good, non-controversial reasons for adopting a “precautionary approach” with regard to possible outcomes such as an epidemic of childhood cancer, and quite what is associated with adopting a precautionary approach to such threats of harm.

“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.”¹ In recent years, public health policy as well as environmental policy have appealed to this, the precautionary principle, by broadening “environmental degradation” to include public health problems.² Despite the effect that “precautionary thinking” has on UK public health policy, most notably in Lord Turner’s report (which recommended on precautionary grounds that mobile phone masts should not be placed near schools), opponents of the principle have claimed that it is, at best, impractical, and, at worst, positively self-contradictory.³ This study outlines a possible basis for the principle—a distinctive account of the structure of harms and benefits—and a possible interpretation of the principle as recommending a particular epistemological strategy in the policy arena. My aim is not to establish definitively that the principle should guide public health policy, but, more modestly, to show that the principle is consistent, defensible and intuitively plausible in certain situations.

I start by assuming two necessary conditions for any philosophical defence of some form of government policy, and thus, for any defence of the precautionary principle, policy proposals should be practical and publicly justifiable.

THREE KINDS OF HARM AND PUBLICITY

The precautionary principle may seem to suggest a strong distinction between the harmful and beneficial outcomes of action. However, this is slightly misleading, as the principle is actually framed in terms of risks of “serious or irreversible damage”, rather than in terms of “any damage”. Therefore, a more accurate interpretation of the precautionary principle is that we (rightly) possess a tripartite division of the outcomes of action. Unlike “normal bad outcomes” (ugly mobile-phone masts), which might be outweighed by “good outcomes” (greater ease of communication), we ought to treat certain outcomes (the avoidable death of innocent children) as “special bad outcomes”, which are not subject to the standard tools of risk–cost–benefit analysis. Are there non-controversial reasons to think that certain bad outcomes ought to be treated as “special”?

Two kinds of argument are presented for such a claim: consequentialist and deontological arguments. I use “consequentialist” here to mean ethical theories which assess actions in terms of the states-of-affairs produced by those actions, and which assume that there is commensurability between different states-of-affairs.⁶ It looks as though it is precisely the assumption of commensurability underlying risk–cost–benefit analysis to which proponents of the precautionary principle object. However, it is possible to argue from a

Correspondence to:
S D John, Department of
History and Philosophy of
Science, University of
Cambridge, Free School
Lane, Cambridge CB2 3RH,
UK; sdj22@cam.ac.uk

Received 15 December 2005
Revised 15 December 2005
Accepted 9 May 2006

consequentialist standpoint, which views all outcomes as ultimately commensurable, to a view that certain outcomes ought to be treated with particular care if we can show that a policy which does not treat those outcomes as special is likely to lead, over time, to worse outcomes than would have occurred had we adopted a policy which treated that class of outcomes as special. We might, as it were, treat certain outcomes as outside the purview of standard risk–cost–benefit analysis not because those outcomes are “really” special, but because we have good long-term reason to treat those outcomes as special.

Such an argument might, in turn, be motivated by inductive evidence. Some of the worst public health crises of the 20th century would have been avoided had we treated the threat of certain “serious or irreversible damages” as special, and, therefore, we ought now to treat such threats as special.⁷ Perhaps the most famous examples of such catastrophes born of misplaced certainty are dichlorodiphenyltrichloroethane in the USA and thalidomide in the UK. So, we need not deny that the respective outcomes, deaths of children from cancer and the increased ease of communication achieved using mobile phones, are commensurable, but we may choose to treat such outcomes as incommensurable (or, at least, as demanding different levels of attention) because of the long-term advantages of such a strategy. This form of argument seems to be in line with one repeated argument for adoption of the precautionary principle in an environmental context, where inductive evidence of past failures of science policy is used to cast doubt on the tools of contemporary risk analysis.⁸

Does a consequentialist argument meet the “publicity condition”? Defenders of precautionary reasoning might respond to this question in a hypothetical mode: if you think that risk–cost–benefit analysis, the method to which precautionary reasoning is normally opposed, meets the publicity condition, then precautionary reasoning is grounded on risk–cost–benefit analysis, and, as such, is as likely as risk–cost–benefit analysis to meet the publicity constraint. The two, usually opposed, approaches, stand or fall together. That there are examples of past cases where scientific risk analysis did go astray is undeniable. However, the question of whether there is enough evidence of such past failures and that past failures can be shown to lead to the systematic conclusion that risk analysis is mistaken is beyond the scope of this paper.

Even if the “consequentialist argument” for the precautionary principle does not succeed, there might be a second way to defend the idea that certain outcomes ought to be treated as “special”. We could argue that government policy ought to be guided not only by considerations of maximising utility but also by recognition of certain institutional obligations. In particular, we might argue that governments have an obligation, above all else, to avoid doing harm to the population, and that this obligation has priority over a weaker positive obligation to do good for the population. O’Neill⁹ has outlined a view of this sort in the global context. If we think that it makes sense to speak of institutional obligations in this way, then a certain set of outcomes are particularly relevant to policy decisions not because of the magnitude of those outcomes, but because for such outcomes to come about as a result of government policy would constitute a breach of the strict negative obligation of non-maleficence. We might resist the thought that all outcomes are commensurable on the grounds that allowing certain outcomes to come about as a result of government policy would be not merely bad, but wrong.

One advantage to be gained from conceptualising the precautionary principle in these “deontological” terms is that it illuminates what Sunstein¹⁰ has attacked as the “conservatism” of precautionary thought—its apparent tendency to block policies that seem likely to have good effects on

the grounds that such policies might also have bad effects. Such conservatism might not be a form of indifference or myopia, but, rather, an expression of a certain conception of the extent of (and limits to) our obligations. Of course, philosophies which emphasise negative obligations seem to rest on a perhaps indefensible distinction between acts of omission and acts of commission; yet framing the precautionary principle in terms of such a fear makes sense, and fits well with an intuitive distinction between positive and negative obligations, as well as the thought that governments, as well as individuals, have such obligations.

However, while it is one thing to say that we often operate with a distinction between different kinds of obligations, it is another to say that this distinction is conceptually stable, and another thing again to say that the distinction is not only stable but capable of forming a “free-standing module” on which there may be an “overlapping consensus”. How, then, should we justify the claim that government policy ought to be guided by certain lexically ordered principles such that we ought to treat certain outcomes as “special”? I do not have space in this paper to argue for such a claim, and if the first “consequentialist” argument discussed earlier is sound, then such an argument might be unnecessary. However, it is worth pointing out that those who argue for a “deontological” account of government policy might have the resources to mount a surprising “indirect” argument for their conclusion, if we take seriously recent research suggesting that “lay” reasoning about risk is not, as risk experts have often claimed, straightforwardly irrational, but expresses a different kind of rationality, which emphasises the binding importance of certain sorts of obligations. Arguments to the effect that lay reasoning about risk is not irrational, but in fact grounded in what we might call a deontological conception of policy, have been mounted in a range of different public health policy debates, most notably in debates over the BSE crisis, and in debates over health and safety, such as railway safety.^{11–12} If the facts of lay reasoning about risk can, indeed, be interpreted in terms of an implicit lay commitment to deontological modes of thought, then it might be possible to put forward a convincing argument to the effect that, even if there are deep philosophical problems with the doing/allowing distinction, if there were strong evidence that enough people believe that the state ought to be guided by deontological considerations, there are “second-order” reasons to structure state policy in these terms. It may be that the demand of neutrality, which seems to suggest that we frame policy decisions in terms of individuals’ preferences, also demands that we take seriously not merely people’s preferences, but also their commitment to certain deontological principles.

PRECAUTION IN PRACTICE: QUESTIONING THE ETHICAL OR EPISTEMIC DIVIDE

Arguably, the precautionary principle relies on a tripartite division of outcomes of action for the purposes of policy making, and there are at least two plausible ways in which we might argue for the adoption of such a tripartite distinction. Therefore, we have good *prima facie* reason to suppose that the underlying assumption of the precautionary principle, that certain sorts of harm are—or ought to be treated as—special, meets the publicity constraint. However, I have suggested that a second constraint on any account of the precautionary principle is that it ought to show us how and when we are to adopt the principle. This “practicality” constraint is particularly important as it has seemed to many that the precautionary principle is too vague to be of any real use in policy making.¹⁰ We need some account of just when we ought to take threats seriously, in the absence of scientific certainty, if we are to generate reasonable

policies. There is, for example, a completely unproven hypothetical threat that playing fields at school cause cancer. No one thinks that there is such a threat, but, once we have removed the demand for scientific certainty from an account of the justification of policy, on what grounds can we justify not taking this hypothetical threat seriously while taking seriously the equally unproven link between mobile-phone masts and childhood cancer? I shall now consider a response to this challenge.

Proponents of the precautionary principle need not deny the legitimacy of risk-cost-benefit analysis for the purposes of policy making. However, we can interpret the precautionary principle as a second-order rule about how we ought to generate factual claims for the purposes of policy making. When engaged in policy making, we often include as facts claims that have been established by statistical testing. We typically design statistical tests to minimise our chance of “false positives”, thus increasing the risk of generating “false negatives”. How are these facts relevant to policy making? Building on Cranor’s¹³ work on the regulation of toxic substances, I suggest the following argument. One way of interpreting the precautionary principle is that we should retain such truth-tropic testing procedures when we are establishing the links between some course of action and “non-special” (good or bad) outcomes for the purposes of policy. However, when we are considering the links between some course of action and “special bad outcomes”, we ought to reverse the burden of proof and instead adopt testing methods that minimise our chance of false negatives, even at the cost of generating more false positives. The fact-like claims derived from these asymmetrical testing procedures should be used as the basis for a risk-cost-benefit calculation. However, if claims such as my hypothetical link between playing fields and cancer cannot be “minimally proven”, even using a lower standard of evidence, then we are justified in not considering the possibility of such outcomes, even on “precautionary grounds”, in the formulation of policy. Hence the popular slogan that precautionary thought includes “reversing the burden of proof”.⁷

The general thrust of my suggested interpretation of the precautionary principle, then, is as a second-order principle governing the generation of claims for inclusion in responsible policy making. Of course, current policy does not use such an “asymmetric” approach to the evaluation and assessment of evidence, and changing policy would be difficult. However, even before we might argue for such a whole-scale restructuring of policy, the suggestion here requires further elaboration and defence. In the remainder of this paper, then, I argue that the claim that we ought to restructure our epistemic policy in line with our ethical goals is not counterintuitive, and I defend my proposal against the charge that it is unscientific or antiscientific.

Philosophers have often treated the normative disciplines of epistemology and ethics as separate domains. However, my proposal above suggests that we might have good “ethical” reasons to adopt particular sorts of “epistemological” strategies in certain circumstances. Is such a suggestion plausible? There is a deep and complex literature, which argues that it is rational to adjust our epistemic strategies in light of the possible pay-offs of being right and wrong. Consider, for example, a situation where you live in an environment that has tigers. You might decide never to form the belief that there is a tiger nearby unless you are absolutely certain that there is one nearby. That is, you might adopt an epistemic policy that favours avoiding false positives. However, there is an obvious downside to this policy: you are also likely to generate false negatives—beliefs that there are no tigers about, when in fact there are tigers.¹⁴ In this case, such beliefs are likely to prove fatal.

Therefore, considering the possible costs of being wrong, it might be rational to adopt a policy where you form the belief that there is a tiger in your environment, and act accordingly even if, in many cases that belief is likely to be false. The alternatives are too unpleasant to bear, even at the cost of believing falsehoods.

Of course, this is an argument about the relationship between prudential ends and belief formation, and, in the form presented here, it assumes the possibility of “epistemic voluntarism”. However, it is not difficult, I suggest, to suppose that many of our belief-generation mechanisms are, in fact, mechanisms that we may choose to adopt or not to adopt, and thus justify the possibility of one form of “epistemic voluntarism”. Furthermore, it is not difficult to see that in a range of social policies, we do allow certain ethical concerns to determine the standard of evidence we demand for forming beliefs. Perhaps the most striking example of such ethical-epistemic interaction is in the courtroom, where our insistence that the innocent man should not be punished determines a procedure which, in many cases, may lead to an incorrect verdict of not guilty.¹⁵ Therefore, I suggest that the general proposal that our ethical and epistemic norms might interrelate in complex ways, particularly in the formulation of policy, is neither surprising, nor obviously confused.

There are two possible responses to the argument I have made here. Firstly, there might be objections that I have misrepresented the nature of statistical testing. When we fail to prove an alternative hypothesis, we do not then say that the null hypothesis has been proven. Rather, in the terminology of statistical testing, we simply say that the null hypothesis has not been disproven. Therefore, it is misleading to suggest that in adopting particular statistical techniques we run the risk of believing falsehoods, thus leading to dreadful consequences. It is, of course, strictly true that when we have not proven the alternative hypothesis, we merely say that the null hypothesis has not been disproven (rather than saying that we believe the null hypothesis). Yet, for all intents and purposes, failure to prove the alternative hypothesis involves acting as if the null hypothesis were true. To phrase the point slightly differently, even if, aware of the subtleties of the philosophy of statistics, we insist that we have not proven the null hypothesis when we return a negative result, the negative result leads to a pattern of action that is indistinguishable from believing the null hypothesis to be true.

Secondly, a rather broader objection to the proposed scheme is that such a scheme would, in some sense, be unscientific, or even antiscientific.¹⁶ I think that it is correct to say that the precautionary principle in general, and my proposed interpretation of the principle in particular, is opposed to the values of science. However, this is only a serious problem if we assume that to adopt any value other than the scientists’ goal of epistemic certainty as the basis for belief formation is inherently irrational. I suggest that, although epistemic caution (only being willing to say that some claim has been proven if one is certain that the claim has been proven) may be one important value, there is no reason to suppose that it is the only value that should regulate epistemic activities. In particular, when we are participating in epistemic activities, which are intended to deliver results for the purposes of policy making, it might be rational to adopt different epistemic goals, in particular the goal of avoiding egregious falsehoods. To adopt such a goal is, in one sense, to deny a fundamental value of science. However, denial of this fundamental value of science is not, itself, necessarily *a priori* irrational, as there may be other values by which epistemic endeavours can be regulated. I have argued that if we take the avoidance of particular outcomes seriously, then avoiding those outcomes is an example of just such a value.

I do not pretend to have shown conclusively that the precautionary principle ought to guide policy; nor do I pretend to have shown that any particular policy that has been justified by appeal to the principle is, in fact, justified. However, I do conclude that the precautionary principle is, properly interpreted, capable of playing an important part in public health policy, in a way that might incorporate our deontological concerns into the powerful framework of risk–cost–benefit analysis. Of course, the conclusion that we ought to rethink the worth of the epistemic values that underlie standard statistical science is likely to seem both disquieting and utopian. However, if we take ethics seriously, then we may have to downplay the value of certainty.

Competing interests: None.

REFERENCES

- 1 **United Nations (UN) General Assembly.** *Rio Declaration on Environment and Development. Report of the United Nations Conference on Environment and Development, Rio de Janeiro, 3–14 June 1992. A/CONF.151/26. Vol I.* New York: UN, 1992.
- 2 **Martuzzi M, Tickner J, eds.** *The precautionary principle: protecting public health, the environment and the future of our children.* Copenhagen: World Health Organization, 2004.
- 3 **Burgess A.** *Cellular phones, public fears, and a culture of precaution.* Cambridge: Cambridge University Press, 2004.
- 4 **Rawls J.** *Political liberalism.* New York: Columbia University Press, 1992.
- 5 **Sunstein C.** *Risk and reason.* Cambridge: Cambridge University Press, 2002.
- 6 **Sen AK, Williams B,** Introduction in: Sen, AK, Williams, B, eds. *Utilitarianism and beyond.* Cambridge: Cambridge University Press, 1982.
- 7 **Harremoës J, Gee D, Macgarvin M, et al.** *The precautionary principle in the twentieth century.* London: Earthscan, 2002.
- 8 **O’Riordan T, Cameron J, eds.** *Interpreting the precautionary principle.* London: Earthscan, 1994.
- 9 **O’Neill O.** *Faces of hunger.* London: Allen and Unwin, 1986.
- 10 **Sunstein C.** *Laws of fear: beyond the precautionary principle.* Cambridge: Cambridge University Press, 2005: Chapter 1.
- 11 **Irwin A.** *Citizen science.* London: Routledge, 1995.
- 12 **Wolff J.** *Railway safety and the ethics of the tolerability of risk.* London: Railway Safety and Standards Board, 2005.
- 13 **Cranor C.** *Regulating toxic substances: a philosophy of science and the law.* Oxford: Oxford University Press, 1997.
- 14 **Godfrey-Smith P,** Signal, decision, action. *J Philos* 1991;**88**:709–22.
- 15 **Hacking, I.** *The taming of chance.* Cambridge: Cambridge University Press, 1990: Chapter 11.
- 16 **Resnik DB.** Is the precautionary principle unscientific? *Stud Hist Philos Biol Biomed Sci* 2003;**34**:329–44.