

GENERAL ETHICS

Should the precautionary principle guide our actions or our beliefs?

M Peterson

J Med Ethics 2007;33:5–10. doi: 10.1136/jme.2005.015495

Two interpretations of the precautionary principle are considered. According to the normative (action-guiding) interpretation, the precautionary principle should be characterised in terms of what it urges doctors and other decision makers to do. According to the epistemic (belief-guiding) interpretation, the precautionary principle should be characterised in terms of what it urges us to believe. This paper recommends against the use of the precautionary principle as a decision rule in medical decision making, based on an impossibility theorem presented in Peterson (2005). However, the main point of the paper is an argument to the effect that decision theoretical problems associated with the precautionary principle can be overcome by paying greater attention to its epistemic dimension. Three epistemic principles inherent in a precautionary approach to medical risk analysis are characterised and defended.

in diagnosing this disease correctly in children give rise to a misdiagnosis rate close to 40%.⁵ However, despite the epistemic uncertainty of the diagnosis, and owing to the possibility of a fatal outcome, the surgeon on duty decides to remove the appendix by laparoscopy as a precautionary measure.

Example 2: Dipyron (Novalgin, metamizole) is a widely prescribed analgesic in South America, Africa, the Middle East and some European countries. In 1973, an estimate suggested an incidence of agranulocytosis of 1 in 3000 patients using dipyron.⁶ On the basis of that estimate, the Swedish MPA forced the producer to withdraw dipyron from the market in 1974. However, the 1973 estimate was soon criticised for having severe methodological flaws. In 1986, the International Agranulocytosis and Aplastic Anemia Study concluded that the risk for agranulocytosis was much lower than believed previously—only 1.1 cases/million users.⁷ Hence in 1995, dipyron was reapproved by the Swedish MPA and prescribed to a limited number of patients. The epistemic uncertainty about dipyron remained however, and after having received 14 new reports of adverse drug reactions the Swedish MPA decided to interdict dipyron again in 1999.⁷ No other drug has ever been interdicted twice in Sweden.

The precautionary principle was originally invoked by policy makers for dealing with environmental issues such as global warming, toxic waste disposal and marine pollution. In recent years, it has also been suggested that the precautionary principle may as well be applied to medical issues. David B Resnik asserts that “properly understood, the [Precautionary Principle] can provide physicians and patients with a useful approach to medical decision making.” (Resnik, p283).¹ See also Resnik,² Weed³ and Alban.⁴

This paper seeks to clarify what the precautionary principle may reasonably be taken to mean in a medical context. As suggested by the title, two different interpretations are considered. According to the normative (action-guiding) interpretation, the precautionary principle should be characterised in terms of what it urges doctors and other decision makers to do. According to the epistemic (belief-guiding) interpretation, the principle should rather be characterised in terms of what it urges us to believe. The difference between the two interpretations is illustrated in the examples below.

Example 1: A 5-year-old girl is brought to the emergency ward by her anxious parents. She has severe abdominal pain, the location of which has not changed since onset. A possible diagnosis would be appendicitis, but difficulties

Examples 1 and 2 illustrate two decisions taken in light of epistemic uncertainty. However, example 1 illustrates a straightforward application of the precautionary principle as a decision rule. The surgeon had to choose between two alternative acts and, naturally, opted for the alternative least likely to cause a fatal outcome. In example 2, the Swedish MPA primarily had to choose what to believe about dipyron. According to the official view, new reliable information had been received after 1995, indicating an unexpectedly high incidence of adverse drug reactions. Therefore, in 1999, the officials of the Swedish MPA decided to believe that the results of the International Agranulocytosis and Aplastic Anemia study were not relevant for the Swedish population (perhaps there was some unknown genetic difference between Swedes and other ethnic groups). The regulatory decision was entirely determined by the epistemic decision. In a certain sense, the officials could not have acted differently once their beliefs had been fixed.

Correspondence to:
Dr M Peterson, Department
of History and Philosophy of
Science, University of
Cambridge, Free School
Lane, Cambridge CB2 3RH,
UK; mbp24@cam.ac.uk

Received 8 December 2005
Revised 11 January 2006
Accepted 13 January 2006

Ethical aspects on epistemic issues are particularly important in medical risk analysis. In situations involving risk and uncertainty, it is far from clear that doctors and other medical decision makers should decide to believe what is most likely to be true.¹ The precautionary principle suggests that doctors should rather seek to acquire beliefs that are likely to protect the patient.

However, despite the recent enthusiasm for applying the precautionary principle in medical risk analysis, several scholars have raised critical concerns about this principle. A common criticism is that the precautionary principle is either too imprecise or tends to be too absolutistic, in that it prohibits activities that, intuitively, ought to be permitted.^{8–11} The two arguments are closely interconnected. The less precise the principle, the more intuitively attractive it seems to become, and vice versa. I state that the problems of using the precautionary principle as a decision rule are indeed insuperable. In fact, there is reason to believe that no version of the precautionary principle can be reasonably applied to decisions that may lead to fatal outcomes, as shown in the next section. However, problems with using the precautionary principle as a decision rule can be overcome by paying greater attention to its epistemic dimension. More precisely put, it can be said that justifiable recommendations about precautionary choices may be derived from epistemic considerations, in conjunction with decision rules that do not assign any particular weight to precaution. Hence, the intuition that medical decision makers ought to be risk averse can be accounted for without applying risk adverse decision rules.

The remaining part of the paper is structured as follows. The next section summarises some major concerns about using the precautionary principle as a decision rule. In the section following that, three epistemic principles inherent in a precautionary appraisal are proposed and briefly explained. The last three sections analyse the three principles in more detail.

THE PRECAUTIONARY PRINCIPLE—SOME PROBLEMS

The proposed contrast between epistemic and normative issues in medical risk analysis can be further clarified by adopting Andrew Stirling's distinction between, on the one hand, precautionary approaches to the regulatory appraisal of risk and, on the other hand, the precautionary principle conceived of as a criterion for choosing among alternative actions.^{12–13} According to this distinction, the precautionary principle is a decision rule that doctors and other decision makers should adopt when making decisions, presumably while taking into account information extracted from risk assessments. A precautionary appraisal of risk is, however, broader. It is a claim about how information has to be processed and analysed on the basis of an appeal to certain epistemic goals. In what follows, the precautionary principle is used for referring exclusively to the normative dimension of the precautionary approach—that is, to a decision rule. The term precautionary appraisal is also given a comparatively narrow meaning, which is likely to differ from Stirling's use of the term. In this paper, this term is used for referring to a set of epistemic principles closely associated with the precautionary principle, mentioned below.

Arguably, there is no such thing as the true formulation of the precautionary principle. There simply exists no single, true formulation of this principle. However, a common point of

¹ Some philosophers would perhaps question the assumption that one can in a genuine sense decide what to believe—beliefs are involuntary. I am aware of that discussion, but will not comment on it here.

departure in many discussions of the precautionary principle is the formulation stated in the Rio declaration:

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. (UNCED, 1993)

This formulation has been extensively discussed in the literature. The aim of this article is not to contribute to this discussion. However, an influential criticism of nearly all formulations of the precautionary principle is that they are, in the words of Daniel Bodansky, “too vague to serve as a regulatory standard” (Bodansky, p 5).⁸ Gray and Bewers¹⁴ develop this criticism further and say that the precautionary principle “poses a number of fundamental problems”, as its logic is unclear and key terms are left undefined. Another line of criticism is that the precautionary principle is absolutist or “overly rigid”.¹⁵ Nollkaemper¹⁶ explicitly raised this concern as he noted that “in several treaties, the precautionary principle is formulated in absolutist terms. It stipulates that once a risk of a certain magnitude is identified, preventive measures to erase that risk are mandatory”. Hence, as virtually every activity is associated with some risk of non-significant damage, it seems that the precautionary principle can therefore be used to prohibit every human activity.^{9–10–17}

In Peterson,¹¹ the problems with using the precautionary principle as a decision rule were articulated by stating an impossibility theorem, showing that no version of the precautionary principle can be reasonably applied to decisions that may lead to fatal outcomes. This result is briefly summarised below.

Consider conditions P, D, A and TO stated in box 1. Condition P is intended to be a partial and very weak formulation of the precautionary principle, which advocates all different versions that this principle ought to agree upon. Condition P can be

Box 1 Principles used for rational decision making

- Precaution (P)
 - If one act is more likely to give rise to a fatal outcome than another, then the second act should be preferred to the first one, given that both fatal outcomes are equally undesirable.
- Dominance (D)
 - If one act yields at least as good outcomes as another under all possible states of the world, and strictly better under some states, then the first act is preferred to the second one.
- Archimedes (A)
 - If the relative likelihood of a non-fatal outcome is increased in relation to a strictly better non-fatal outcome, then there is some decrease of the relative likelihood of a fatal outcome that counterbalances this precisely.
- Total order (TO)
 - Preferences between acts are complete, antisymmetric and transitive.

conceived of as a minimal condition that ought to be implied by plausible versions of the precautionary principle. The three other conditions—D, A and TO—are general normative conditions that every principle used for rational decision making ought to satisfy.

The following example illustrates the reasoning behind the dominance principle (D). Suppose that your doctor has instructed you to take pill Q if you start feeling ill. You now start feeling ill. You have been informed that taking the pill will not give rise to any adverse drug reactions. Further, as you have already bought the pill and it cannot be stored until you feel ill next time, it would not cost you anything to take it. Finally, the pill does not taste bad; it tastes of strawberry. Given this information, should you take the pill or not? The point is that no matter whether the pill actually cures your disease, you will do at least as well if you take the pill as if you do not. Therefore, according to the dominance principle, you should take the pill. You simply have nothing to lose.

The Archimedean condition articulates the plausible intuition that, everything else being equal, if the relative likelihood of a non-fatal outcome (say, gastric ulcer) is increased in relation to a strictly better non-fatal outcome (say, a headache), the act thereby becomes slightly worse. This can, however, be counter-balanced by decreasing the relative likelihood of a fatal outcome (death), thereby improving the act equally much.

Total order is a technical condition. A preference ordering is complete if and only if, for every pair of acts X and Y, act X is at least as preferred as Y, or Y is at least as preferred as X. Antisymmetry means that if X is at least as preferred as Y and Y is at least as preferred as X, then X and Y are equipreferred. Finally, if a preference ordering is transitive it means that if X is preferred to Y, and Y is preferred to Z, then X is preferred to Z.

The four conditions stated above—P, D, A and TO—are logically inconsistent.¹¹ A summary of the proof is given in the appendix. Arguably, this impossibility theorem shows that the “cost” of accepting the precautionary principle as a decision rule is too high, as condition P cannot be accepted unless at least one of the three other conditions is given up. This indicates that intuitions about precaution should not be explicated in normative terms.

In the remaining sections of this article, the epistemic dimension of the precautionary concept is analysed.

THREE EPISTEMIC PRINCIPLES

The epistemic principles relevant in medical risk analysis can be characterised by adopting the widespread definition of risk analysis as a process consisting of three phases—that is, hazard identification, risk assessment and risk management. Given this trichotomy, a precautionary appraisal could be characterised as (an epistemic) recommendation about the proper way to identify hazards and assess risks; not as (a normative) claim about proper risk management.

What epistemic principles would be inherent in a precautionary appraisal of medical risks? A common suggestion, which has been much discussed in the literature, is that in risk assessments it is more desirable to avoid false negative errors than false positive ones.^{10 18 19} If this epistemic principle, the preference for false positives, is valid, it would be more undesirable from an epistemic point of view to not discover a relationship between a hazard and an activity that is in fact there, compared with incorrectly discovering a relationship that is actually non-existent.^{19 20} This is usually not thought to be the case in scientific research, as scientists prefer to remain unknowing about a truth, rather than believing something that is actually false. If the preference for false positives is to be accepted, it must be justified in some other way.

I propose that the second epistemic principle inherent in a precautionary risk appraisal should be the ecumenical principle. This principle holds that in case the experts’ views on some risk issue conflict, they are allowed to adopt (and act from) any of those alternative views. Or, put in other words, the views of all sufficiently qualified experts should be regarded as legitimate, not only those of the most prominent expert.

The third epistemic principle I would like to suggest is the principle of non-monotonicity. This principle denies that “more is always better” when it comes to the amount of (relevant) information included in a risk assessment. A noteworthy implication of this principle is that there are cases in which information should be excluded from a risk assessment, even though the information would be relevant. Briefly put, the motivation is that a decision maker faced with too much information might be unable to see the wood for the trees. Therefore, the epistemic value of a risk assessment does not strictly increase with respect to the amount of information contained in it.

The three epistemic principles are summarised below.

- *The preference for false positives:* In a precautionary appraisal of medical risks, it is more desirable from an epistemic point of view to avoid false negative errors compared with false positive ones.
- *The ecumenical principle:* In a precautionary appraisal of medical risks, all expert views should be regarded as legitimate, not only the view put forward by the most prominent or influential expert.
- *The principle of non-monotonicity:* In a precautionary appraisal of medical risks, the epistemic value of a risk assessment is not strictly increasing with respect to the amount of information contained in it.

The ecumenical principle and the principle of non-monotonicity sometimes yield conflicting epistemic recommendations. According to the ecumenical principle, all expert views should be considered seriously. However, according to the principle of non-monotonicity, a decision maker could be faced with too much information. The solution to this problem is to conceive of all the epistemic principles as *prima facie* principles.²¹ All principles are indeed valid, but they might be overridden by other considerations. The epistemic duty proper (ie, what one ultimately ought to believe) can be determined only after considering all relevant principles. Such an overall judgement about the epistemic situation typically involves trade-offs between conflicting considerations.

THE PREFERENCE FOR FALSE POSITIVES

The first epistemic principle holds that in a medical risk appraisal, it is more desirable to avoid false negative errors than false positive ones. This might seem like an unintuitive principle. After all, it is exactly the other way around in the sciences, so why should a risk appraisal be any different? Arguably, the answer is that the aim of science differs from that of medical risk appraisals. Scientists strive to acquire as many true beliefs as possible, while minimising the false ones. However, the aim of a medical risk appraisal is not to provide a correct representation of medical facts. The aim is rather to protect patients from medical hazards.

If offered a choice between failing to reject a hypothesis that is in fact false and failing to adopt a hypothesis that is in fact true, scientists would generally prefer to not discover an additional truth about the world compared with coming to believe something that is in fact false. A simple and sound explanation of this epistemic preference is seen. New scientific beliefs are often instrumental when making further discoveries, so any mistake incorporated into the corpus of scientific

knowledge is likely to give rise to more mistakes further down the road. This is illustrated by the well-known example of phlogiston. In the 17th and 18th centuries, it was widely accepted that all flammable materials contained phlogiston, a substance claimed to have no mass, colour, taste or odour. It was believed that phlogiston was given off in combustion. This false belief guided chemists in the wrong direction for a long period. In fact, chemists did not come any closer to the truth about combustion for more than a century. The mistake of believing in phlogiston was not corrected until 1777, when Lavoisier²² presented his theory of combustion. So, briefly put, the scientists' preference for false negatives can be traced to the negative consequences for future research of incorrectly accepting a false hypothesis.

What about medical risk appraisals? The most plausible argument for preferring false positive errors over false negatives is, arguably, that the consequences of coming to believe that something is hazardous when in fact it isn't are seldom disastrous. The consequences of falsely believing something to be safe when it isn't might, however, be disastrous. If I believe that it is safe to drink the tap water when it isn't, I might get sick. Hence, it is better to pay a small amount for a bottle of mineral water. Call this the argument from decision theory.

The argument from decision theory relies on several empirical premises. These can be articulated by tackling the following problem suggested by Tim Lewens (personal communication): You live in a jungle populated by an unknown number of tigers. The tigers are yellow and black. Unfortunately, everything eatable in the jungle is also yellow. For example, bananas are yellow. You decide to protect yourself against tigers by building a device that detects and warns for everything that is yellow. The good news is that because of the detector you will not be killed by a tiger. The bad news is that you will starve to death, because you will never find anything to eat. Hence, it is far from clear that it is in general better to prefer false positives over false negatives.

The tiger example makes it clear that the epistemic preference for false positives would only be acceptable if we had reasons to believe that the combined undesirability and likelihood of making a false positive error outweighs the combined undesirability and likelihood of making a false negative error. Proponents of the argument from decision theory believe that we have such reasons. Of course, in the tiger example, the number of tigers in the jungle might be very small, whereas the consequence of not finding any bananas to eat might be disastrous. Under these circumstances, a preference for false positives would be unreasonable. However, in many real-life situations, there are empirical reasons indicating that the risk of missing a real hazard outweighs the consequence of making a false negative error. Metaphorically speaking, this means that the number of tigers is so high that it outweighs the fact that no bananas are found. At least in a one-shot decision—that is, a decision that is never repeated, this could motivate the principle of preferring false positives.

The principle of preferring false positive errors is often combined with the claim that the burden of proof should be reversed when risks are high.²³ According to this view, it is not the person who claims that X is hazardous who has the burden of proof; it is rather the person who claims that X is safe who ought to support his claim with arguments. This idea about a reversed burden of proof is, however, problematic. Arguably, anyone who is making a claim about something has the burden of proof, no matter what the claim is. To see this, suppose that there exists a set of beliefs B, such that one is free to accept these beliefs without having any reason for doing so—that is, without having any burden of proof. Let *b* be an element of B.

Then consider a person who happens to believe *not-b*, and does so for some reason. For example, let *not-b* be the belief that a new drug does not give rise to any adverse drug reactions; the reason might be that preliminary, inconclusive tests give partial support to this belief. Now, faced with the belief *b*, the agent has to decide whether to revise her previous belief, *not-b*, or reject the new belief *b*. As *not-b* and *b* are contradictory, both beliefs cannot be accepted. However, if the claim about a fixed burden of proof is taken seriously, it would imply that a person who believes *not-b* for some reason, which might be inconclusive, would be forced to give up that belief in favour of the opposite belief *b*, without being able to give any reason for this revision of beliefs. This is implausible. In fact, it is almost bizarre to accept a principle forcing us to change beliefs without being able to give any reason for doing so.

At this point, it might be objected that the idea of a reversed burden of proof is applicable only to cases in which one has not yet acquired a belief in either *not-b* or *b*. Claims about a reversed burden of proof can, therefore, be invoked only if it is completely open whether one should believe *not-b* or *b*. Given this qualification, the problem outlined above could be avoided. Unfortunately, the qualification also makes the claim more or less empty. In nearly every case of practical relevance, people already hold some belief about the issue under consideration. Consider, for example, the case of genetically modified food. If the claim about a reversed burden of proof is taken seriously, one should believe that genetically modified food is hazardous until it has been proven safe. The problem is, however, that most people already hold a belief about genetically modified food, and some people do indeed believe that genetically modified food is safe. Should they really change their view, without being able to give any reason for doing so?

Note that the preference for false positives can be accepted without simultaneously adopting the idea about a reversed burden of proof. The two principles are distinct. The first is a methodological rule derived from statistics, according to which it is less serious, in a risk appraisal, to make false positive errors compared with making a false negative error. The second is a more general metaepistemological principle about how one should decide what to believe.

THE ECUMENICAL PRINCIPLE

It is not uncommon that experts disagree. In the debate on dipyrone, some experts said that it would be appropriate to approve dipyrone a second time, whereas other experts disagreed.²⁴ Both parties had access to the same raw data, but they interpreted the data differently.⁷ In cases where experts disagree, it is often difficult for the decision maker to take this disagreement into account in a reasonable way. In many cases, one simply has to decide which expert appears to be most trustworthy. Arguably, this is something that could be questioned from an epistemic point of view. According to the ecumenical principle, all expert views should be considered in a precautionary appraisal, not only the views put forward by the most prominent or influential expert.

The ecumenical principle can be formulated in more detail by applying the machinery of deontic logic to doxastic states. Instead of asking which opinion is most probable to be true, we may ask what a rational person is permitted to believe is true. More precisely put, for each proposition *x* we assume that it is forbidden (F*x*), or permitted (P*x*), or obligatory (O*x*) to believe that the proposition is true. Of course, (1) everything that is obligatory to believe is also permitted to believe, and (2) everything that is forbidden to believe is not permitted to believe, and (3) everything that is permitted to believe is not forbidden to believe. Let *x* be an arbitrary proposition. Consider the following version of the ecumenical principle.

1. In a precautionary appraisal of medical risks, it is obligatory to believe that *x* if and only if every expert believes that *x*.
2. In a precautionary appraisal of medical risks, it is permitted to believe that *x* if and only if at least some experts believe that *x*.
3. In a precautionary appraisal of medical risks, it is forbidden to believe that *x* if and only if no expert believes that *x*.

The main advantage of adopting a deontic formulation of the ecumenical principle is that the principle then becomes no more precise than what is justified by the experts' judgements. If some quantitative (probabilistic) principle is adopted for reconciling divergent expert opinions, the policy maker may probably be presented with material that appears to be much more precise than it actually is.

The ecumenical principle has several interesting corollaries. First and foremost, if some experts believe that *x*, and some believe that not-*x*, then they are permitted to believe that *x* and permitted to believe that not-*x*. However, it does not follow that they are permitted to believe *x* and not-*x*. In a similar vein, it follows that it is obligatory to believe not-*x* just in case it is forbidden to believe *x*, granted the law of the excluded middle.

The ecumenical principle can be directly applied to the dipyrone example. As some experts believed that the incidence of agranulocytosis linked to dipyrone was high (about 1 in 3000), whereas another group of experts believed that the incidence was low (about 1 in 1 000 000), both views were permitted to be believed. In fact, a study in 2002 reviewed all spontaneous reports from Swedish hospitals between 1965 and 1999 of serious blood dyscrasias associated with dipyrone. The incidence was reported to be 1 case/1431 prescriptions ($n = 66$, of which 52 occurred before 1974 and 14 between April and September 1995).⁷ Hence, it would also be permissible to believe that the incidence was about 1 in 1400. However, given these epistemic facts, regulators subscribing to the precautionary approach might then go on and defend a normative principle prescribing that one should act as though the incidence was high—that is, about 1 in 1400.

It might be objected that the proposed model for risk appraisal makes the most pessimistic experts too influential. According to my model, the influence of one or a few pessimistic experts can never be counterbalanced by any number of significantly more optimistic experts. However, this is also the reason why the model is a precautionary model. Personally, I would therefore be prepared to accept this implication. Nevertheless, if the suggested notion of epistemic precaution is judged to be too extreme, one could strengthen the criterion of doxastic permissibility by requiring more from a proposition that it is permissible to believe in—for example, by requiring that a sufficiently influential or a sufficiently large number of experts have to believe in the proposition in question.

THE PRINCIPLE OF NON-MONOTONICITY

The principle of non-monotonicity holds that “more is not always better”. There are epistemic situations in which decisions will be worse if more information is acquired. This is a controversial claim, and the principle could be easily misinterpreted in a way that would make it trivially false. For example, there is no reason to believe that an ideal decision maker, with unlimited computing capacity, could ever fail to make a decision that is at least as good as before by acquiring more information, provided that no old information is rejected or ignored. The principle of non-monotonicity could also be misinterpreted in a way that would make it trivially true: It is easy to imagine that a non-ideal decision maker, with limited computing capacity, would sometimes make a worse decision

after having acquired more information simply because he failed to process the huge amount of information available to him. None of these interpretations of the principle of non-monotonicity will be given any further consideration here.

According to the interpretation of the principle of non-monotonicity considered in this paper, there are epistemic situations in which decisions will become worse if more information is acquired, and this holds true even if the decision is taken by an ideal decision maker. Imagine a new drug which is to be approved by some regulatory agency. Initial tests suggest that the incidence of some adverse drug reaction, say agranulocytosis, is about 1 in 1 000 000. On the basis of this piece of rather imprecise information, we may assume that the agency would be prepared to approve the new drug, given that (1) it is at least as good as previous substances and (2) the incidence of agranulocytosis and other adverse drug reactions is no higher than for similar substances. However, the regulatory agency then acquires more information. The incidence of agranulocytosis is not randomly distributed in the population. In fact, there is reason to believe that only patients who are bearers of some yet undiscovered gene will contract agranulocytosis when treated with the new drug. On the basis of this enhanced information, the regulatory agency then decides that the new drug can be approved only if the gene causing agranulocytosis is identified. This would allow doctors to make genetic tests before prescribing the drug to patients. Unfortunately, numerous examples indicate that commercial companies requested to provide this type of information very often conclude that the costs of identifying the relevant gene would not exceed the expected profits. Therefore, the gene will never be identified and the new drug will never be approved. This is a pity, as the aggregated amount of human suffering could have been decreased by approving the new drug, even if the relevant gene was not identified, as the new drug was in fact more efficient than the old one.

The agranulocytosis example indicates that in some cases it is better, when making a precautionary risk appraisal, to believe that some hazard is randomly distributed rather than deterministically distributed, given that there is no practically feasible way to find out who will be affected by the hazard. The veil of ignorance surrounding a random distribution helps the decision maker to make better decisions. This holds true even if the decision maker is an ideal person who is able to process unlimited amounts of information in virtually no time.

CONCLUSION

This article has explored a distinction between normative and epistemic issues in a precautionary approach to medical risk analysis. Three epistemic principles have been characterised. Together they determine which “decision matrix” or “problem specification” a normative principles will be applied to. It is important to keep in mind that even if the epistemic principles inherent in a precautionary appraisal are adopted, it does not follow that the decision maker has to decide which action to take by adopting the corresponding normative principle. Nothing prevents the decision maker from formulating his decision problem by adopting the three epistemic principles proposed above, and then choosing what to do by applying a completely risk-neutral normative principle.

ACKNOWLEDGEMENTS

I am especially grateful for the support of David Slavin, and also to Pfizer Global Research and Development. John Cantwell, Richard Jennings, Stephen John, Tim Lewens and Per Sandin have given very helpful comments on earlier drafts.

Competing interests: MP's work on this article has been supported by a generous grant from Pfizer Global Research and Development.

REFERENCES

- 1 Resnik D. The precautionary principle and medical decision making. *J Med Philos* 2004;**29**:281–99.
- 2 Resnik D. Is the precautionary principle unscientific? *Stud Hist Philos Sci Part C. Stud Hist Philos Biol Biomed Sci* 2003;**34**:329–44.
- 3 Weed DL. Precaution, prevention, and public health ethics. *J Med Philos* 2004;**29**:313–32.
- 4 Alban S. The “precautionary principle” as a guide for future drug development. *Eur J Clin Invest* 2005;**35**:33–44.
- 5 Paris CA, West EJ. Abdominal pain in children and the diagnosis of appendicitis. *West J Med* 2002;**176**:104–7.
- 6 Böttiger LE, Westerholm B. Drug-induced blood dyscrasias in Sweden. *BMJ* 1973;**3**:339–43.
- 7 Hedenmalm K, Spigset O. Agranulocytosis and other blood dyscrasias associated with dipyron (metamizole). *Eur J Clin Pharmacol* 2002;**58**:265–74.
- 8 Bodansky D. Scientific uncertainty and the precautionary principle. *Environment* 1991;**33**:4–5, 43–4.
- 9 Manson N. The precautionary principle. The catastrophe argument, and Pascal’s wager. *Ends Means* 1999;**4**:12–16.
- 10 Sandin P. *Better safe than sorry: applying philosophical methods to the debate on risk and the precautionary principle* (diss.) Theses in philosophy from the Royal Institute of Technology, ISBN 91-7283-907-4, 2004.
- 11 Peterson M. The precautionary principle is incoherent. *Risk Analysis* 2006;**26**(3):595–601.
- 12 Stirling A. Science and precaution in the appraisal of electricity supply options. *J Hazard Mat* 2001;**86**:55–75.
- 13 Stirling A, Mayer S. Finding a precautionary approach to technological developments? Lessons for the evaluation of GM crops. *J Agric Environ Ethics* 2002;**15**:57–71.
- 14 Gary JS, Bewers M. Towards a scientific definition of the precautionary principle. *Marine Pollution Bulletin* 1996;**32**(11):768–71.
- 15 Bodansky D. Commentary: the precautionary principle. *Environment* 1992;**34**:2–4.
- 16 Nollkaemper A. The precautionary principle in International Environmental Law: what’s new under the sun? *Marine Pollut Bull* 1991;**22**:107–10.
- 17 McKinney WJ. Prediction and Rolston’s Environmental Ethics: Lessons from the philosophy of Science. *Science and Engineering Ethics* 1996;**2**(4):429–40.
- 18 Holm S, Harris J. Precautionary principle stifles discovery. *Nature* 1999;**400**:398.
- 19 John SD. The will to false negatives: making some sense of the precautionary principle. *J Med Ethics* (in press).
- 20 Harremoës P, Gee D, MacGarvin M, et al. *Late lessons from early warnings: the precautionary principle 1896–2000*. Copenhagen: European Environment Agency, 2002.
- 21 Ross WD. *The right and the good*. New York: Oxford University Press, 1930.
- 22 Lavoisier A. *Traité élémentaire de chimie, présenté dans un ordre nouvell et d’après les découvertes modernes, 2 vols*. Paris: Chez Chuchet, 1789.
- 23 van den Belt H. Debating the precautionary principle: guilty until proven innocent or innocent until proven guilty? *Plant Physiol* 2003;**132**:1122–6.
- 24 Edwards R. Adverse drug reactions: finding the needle in the haystack. *BMJ* 1997;**315**:500.

APPENDIX

PROOF OF IMPOSSIBILITY THEOREM

Let $\{x_1, x_2, \dots\}$ be a set of outcomes produced by a set of states S and a set of acts. Acts are conceived of as vectors of outcomes, $X = [x_1, \dots, x_n]$. The notation is chosen such that the most likely outcome is listed first, and so on. Let \geq_p be a binary relation on outcomes denoting relative likelihood—that is, a measure of qualitative probability. Let \geq_d be a binary relation on outcomes that orders them from the most desirable ones to the least desirable. The letters a, b, c, \dots represent degrees of desirability. Fatal outcomes, such as death, are denoted by the letter f . The notation is chosen such that $a \geq_d b \geq_d \dots \geq_d f$, where each degree of desirability corresponds to a possible outcome x_1, x_2, \dots . The relations approximately_d and >_d are defined in terms of \geq_d in the usual way. The relation $>$, without an index, is a preference relation on the set of acts.

The following conditions correspond to the intuitive formulations of P, D, A and TO stated above.

- P: Let $X = [x_1, \dots, x_n]$ and $Y = [y_1, \dots, y_n]$ such that for exactly one x_i and one y_j , $x_i \sim_d y_j \sim_d f$. Then, if $y_j >_p x_i$, it holds that $X > Y$.
- D: If $x_i \geq_d y_i$ for all i , and there is some j such that $x_j >_d y_j$, then $X > Y$.
- A: Let x^I denote the outcomes produced by the subset I of the states S . Then, for every $X = [a, b, \dots, f, c]$, there are some J, K, L, M such that $[a^J, b^K, \dots, f^L, c^M] \sim [b^J, a^K, \dots, c^L, f^M]$.
- TO: The relation $>$ is complete, asymmetric and transitive.

THEOREM: Conditions P, D, C and TO are logically inconsistent.

PROOF OF THEOREM: Let $X = [a, b, a, f, c]$. Condition A then implies that there are some J, K, L, M such that $X' \sim X''$, where $X' = [a^J, b^K, a, f^L, c^M]$ and $X'' = [b^J, a^K, a, c^L, f^M]$. Let $Y = [b^J, a^K, b, c^L, f^M]$. Then, P implies that (1): $Y > X'$. Furthermore, condition D implies that $X'' > Y$. As $X' \sim X''$, it follows that $X' \sim X'' > Y$. Condition TO guarantees that the preference ordering is transitive, so (2): $X' > Y$. This contradicts (1), as TO also guarantees that the preference ordering is asymmetric.