Foundations of Science (2007) 12:257–268
DOI 10.1007/s10699-007-9107-6

Values in pure and applied science

Sven Ove Hansson

Published online: 28 March 2007

Abstract In pure science, the standard approach to non-epistemic values is to exclude them as far as possible from scientific deliberations. When science is applied to practical decisions, non-epistemic values cannot be excluded. Instead, they have to be combined with (value-deprived) scientific information in a way that leads to practically optimal decisions. A normative model is proposed for the processing of information in both pure and applied science. A general-purpose corpus of scientific knowledge, with high entry requirements, has a central role in this model. Due to its high entry requirements, the information that it contains is sufficiently reliable for the vast majority of practical purposes. However, for some purposes, the corpus needs to be supplemented with additional information, such as scientific indications of danger that do not satisfy the entry requirements for the corpus. The role of non-epistemic values in the evaluation of scientific information should, as far as possible, be limited to determining the level of evidence required for various types of practical decisions.

Keywords Corpus · Values in science · Epistemic values · Non-epistemic values · Scientific values · Applied science · Pure science

1 Introduction

In his seminal paper on values in science, Rudner (1953:2) referred to practical uses of science such as ascertaining that "a toxic ingredient in a drug was not present in lethal quantity". In contrast, the debate on values in science that his paper gave rise to has had a strong emphasis on so-called pure science, i.e. science that aims at knowledge per se rather than knowledge for practical use (Jeffrey 1956; Levi 1962; Hempel 1960;

S. O. Hansson (✉)
Department of Philosophy and the History of Technology
Royal Institute of Technology
Teknikringen 78, 100 44 Stockholm, Sweden
e-mail: soh@kth.se

Feleppa 1981; Harsanyi 1983). This change of emphasis is part of a more general pattern: the philosophy of science has with few exceptions focused almost exclusively on "pure" science, and very little effort has been spent on the philosophical issues that are peculiar to "applied" science. It is the purpose of the present study to clarify how the role of non-epistemic values in applied science differs from their role in pure science.

This undertaking is somewhat complicated by the problematic nature of the traditional division between "pure" and "applied" science. Although this categorization is taken for granted in most discussions on research policy, it is often very difficult to apply it to actual research projects. Many projects have the characteristics of both pure and applied science: they are interesting from the "pure" perspective of satisfying our curiosity about nature, but at the same time they bring promises of practical use. This applies not least in currently large research areas such as biochemistry and materials science. It would be more accurate to treat explanatory (intra-scientific) fruitfulness and practical (extra-scientific) usefulness as two properties that one and the same scientific investigation can have to different degrees, rather than to divide such investigations into the two categories pure and applied (Hansson 2006b).

For the present purpose, however, it is expedient to retain the traditional concepts. In what follows, science will be called "pure" when it is assumed to aim at knowledge that is not primarily valued for its practical usefulness, and "applied" when it is assumed to aim primarily at solving practical problems. It should be remembered that often one and the same scientific activity can be seen in both perspectives.

The purpose of the present contribution is normative, i.e. its ultimate aim is to answer the question how non-epistemic values should be treated in pure, respectively, applied science, in order to make science an as efficient tool as possible for human strivings. However, the distinction between normative and descriptive approaches is not crystal clear in the philosophy of science. I do not profess to have invented entirely new normative principles for science. Instead I will attempt to explicate, systematize, and to some extent reform what can be described as the underlying ethos of science. Hence my attitude to scientific practices with respect to values is similar to the attitude that most medical ethicists have to medical practice.

Values enter science in many ways, for instance in our decisions what to investigate, how to investigate it, and what conclusions to draw from our investigations. The present study is focused on the latter area of value influence. Hence, I will discuss the influence of values on the conclusions that we draw from scientific investigations, not their influence on the planning and conduct of such investigations.

In Sect. 2, a simple model is provided of the knowledge process in pure science. In Sect. 3, alternatives to this model are presented and the model is defended against them. In Sect. 4, the model is extended to cover decision processes in applied science. Alternatives to this extended model are presented in Sect. 5, and again, the model is defended. In Sect. 6, the major conclusions of this investigation are summarized.

2 Values in pure science

Scientific knowledge begins with data that originate in experiments and other observations. Through a process of critical assessment, these data give rise to the scientific corpus (See Fig. 1). The corpus can, roughly, be described as consisting of those statements that could, at the time being, legitimately be made without reservation in a

Fig. 1 The knowledge formation process in pure science

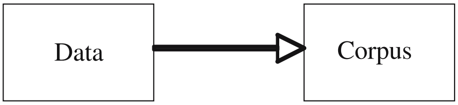

(sufficiently detailed) textbook. Alternatively it can be described as consisting of that which is taken for given by the collective of researchers in their continued research, and thus not questioned unless new data give reason to question it (Hansson 1996).

The corpus consists of generalized statements that describe and explain features of the world we live in, in terms defined by our methods of investigation and the concepts we have developed. Hence, what enters the corpus is not a selection of data but a set of statements of a more general nature. Whereas data refer to what has been observed, statements in the corpus refer to how things are and to what can be observed. Hypotheses are included into the corpus when the data provide sufficient evidence for them, and the same applies to corroborated generalizations that are based on explorative research.[1]

The scientific corpus is a highly complex construction, much too large to be mastered by a single person. Different parts of it are maintained by different groups of scientific experts. These parts are all constantly in development. New statements are added, and old ones removed, in each of the many subdisciplines, and a consolidating process based on contacts and cooperations between interconnected disciplines takes place continuously. In spite of this, the corpus is, at each point in time, reasonably well-defined. In most disciplines it is fairly easy to distinguish those statements that are, for the time being, generally accepted by the relevant experts from those that are contested, under investigation, or rejected. Hence, although the corpus is not perfectly well-defined, its vague margins are fairly narrow.

The process that leads to modifications of the corpus is based on strict standards of evidence that are an essential part of the ethos of science. When determining whether or not a new scientific hypothesis should be accepted for the time being, the onus of proof falls squarely to its adherents. Similarly, those who claim the existence of an as yet unproven phenomenon have the burden of proof. In other words, the corpus has high entry requirements. This is essential to prevent scientific progress from being blocked by wishful thinking and from the pursuit of all sorts of blind alleys. We must be cautious with what we take for granted in our scientific work. But of course there are limits to how high the requirements can be. We cannot leave everything open. We must be prepared to take some risks of being wrong, but these must be relatively small risks.

The entry requirements of the corpus can be described in terms of how we weigh the disadvantages for future research of unnecessarily leaving a question unsettled against those of settling it incorrectly. This is closely related to what values we assign to truth and to avoidance of error. In addition, our decisions on corpus inclusion can be influenced by other values that concern usefulness in future science, such as the simplicity and the explanatory power of a theory. All these are values, but they are not moral values. Hempel called them epistemic utilities and delineated them as follows:

[1] Large parts of modern science are explorative rather than hypothesis-testing. Hence, investigations aimed at determining a DNA sequence or the structure of a molecule are not hypothesis-testing (Hansson 2006a).

> "[T]he utilities should reflect the value or disvalue which the different outcomes have from the point of view of pure scientific research rather than the practical advantages or disadvantages that might result from the application of an accepted hypothesis, according as the latter is true or false. Let me refer to the kind of utilities thus vaguely characterized as *purely scientific*, or *epistemic, utilities*." (Hempel 1960: 465)

Whereas epistemic values determine what we allow into the corpus, influence from non-epistemic values is programmatically excluded. According to the ethos of science, what is included in the corpus should not depend on how we would like things to be but on what we have evidence for. Therefore, it is part of every scientist's training to leave out non-epistemic values from her scientific deliberations as far as possible. This, of course, is not perfectly achieved. As was noted by Ziman, we researchers all have interests and values that we try to promote in our scientific work, "however hard we try to surpass them". But as he also noted, "the essence of the academic ethos is that it defines a culture designed to keep them as far as possible under control" (Ziman 1996: 72).[2]

3 Alternative views

It can hardly be denied that science both includes and has to include a process in which scientific statements are selected for acceptance. However, it does not follow that this process must have the characteristics outlined above. At least two alternative accounts should be considered before we turn to the more complex issues of applied science. One of these approaches is to replace the single corpus by several corpora, perhaps one for each discipline. The other is to replace the sharp limit between elements and non-elements of the corpus by a system with multiple degrees of acceptance, such that most scientific statements have some intermediate degree of acceptance between the highest degree and outright rejection. We can call these alternatives "multiple corpora", respectively, "fuzzy corpus".

3.1 Multiple corpora

Beginning with *multiple corpora*, this idea assumes that the different disciplines are largely independent of each other. Arguably this was the case in the early days of science, but today the interdisciplinary interconnections in science are strong and rapidly strengthening. In the last half century or so, integrative disciplines such as astrophysics, evolutionary biology, biochemistry, ecology, quantum chemistry, the neurosciences, and game theory have developed at dramatic speed and contributed to tying together previously unconnected disciplines. The community of interdependent disciplines includes not only those academic disciplines that are covered by the restrictive English term "science" but also the wider range of disciplines that are covered by the German term "Wissenschaft".

[2] Probably, the largest deviations from this ideal concern non-controversial values, i.e. values that are shared by virtually everyone or by everyone who takes part in a particular discourse. The presence of non-controversial values in the corpus is particularly problematic in cases where the value consensus among experts in a particular field is not shared by everyone else. Consensus views among economists about economic growth have been attacked by non-economists, and feminists have uncovered androcentric values that were uncontroversial in expert communities dominated by men.

Somewhat paradoxically, belief in the coherence of science seems to have been much stronger in the first half of the 20th century than what it is today. Although the reductive account of relations between the disciplines that was popular at that time is not tenable, it is remarkable that interdisciplinary interdependence has increased dramatically in science at the same time as belief in it seems to have receded.

Hence, in actual practice science operates with a common corpus with increasingly strong interconnections. It would of course in principle be possible to claim that nevertheless, normatively, the corpus should be broken up into parts with less interdependence. However, it is difficult to see how, on any reasonable account of the aims of science, such an approach could be justified. The alternative with multiple corpora does not seem plausible.

3.2 A fuzzy corpus

We can therefore turn to the other alternative, that of a fuzzy corpus. In the model proposed in Sect. 2, scientific statements are subject to a binary classification, as accepted (into the corpus) or not accepted.[3] There is an obvious way to replace this classification with many degrees of acceptance: we can use Bayesian decision theory (Jeffrey 1956). According to the Bayesian ideal of rationality, all statements about the world should have a definite probability value assigned to them. Contingent propositions should never be fully believed, but can be assigned high non-unit probabilities. The resulting belief system is a complex web of interconnected probability statements.

There is a prominent feature of actual human belief systems that Bayesian models do not take into account: the cognitive limitations of human beings. These limitations are in fact so severe that in order to arrive at a manageable belief system we have to "fix" a large amount of our beliefs to (provisional) certainty, and take as true (false) much of that to which we would otherwise only assign a high non-unit (low non-zero) probability.[4] In order to grasp complex situations, we reduce the prevailing epistemic uncertainty to full beliefs.[5] The mother fully believes that the child she rears is the child to which she gave birth, in spite of the slight probability that there was an exchange of babies in the maternity ward. The Bayesian mother would only assign a probability close to 1 to this statement. This process of uncertainty-reduction, or "fixation of belief" (Peirce 1934), helps us to achieve a cognitively manageable representation of the world, thus increasing our competence and efficiency as decision-makers.

In the collective processes of science, just as in individual belief formation, our cognitive limitations make such a reduction process necessary. Here, as well, our cognitive limitations make massive reductions of high probabilities to full beliefs (certainty) indispensable if we wish to be capable of reaching conclusions and making decisions. As one example of this, since all measurement practices are theory-laden,

[3] Probabilities have a role in science, but according to this model they appear on another level: the scientific statements about the world that are included in the corpus are often expressed in probabilistic terms.

[4] On the other hand, we also regard many issues as unsettled or uncertain, but do not assign definite probabilities or degrees of belief to them. Thus, whereas a (hypothetical) Bayesian subject assigns probabilities distinct from 0 and 1 to all contingent factual statements, actual subjects have very few such probabilistic beliefs but instead (i) judgments held to be true or false, and (ii) judgments that are unsettled but to which no exact numerical probability has been assigned.

[5] The word "reduce" is used metaphorically. I do not wish to imply that all full beliefs have been preceded by more uncertainty-laden belief states, only that they can be seen as reductions in relation to an idealized belief state in which uncertainty is always fully recognized.

no reasonably simple account of measurement would be available in a Bayesian approach (McLaughlin 1970). The corpus can in fact be seen as the outcome of a collective epistemic reduction process. This process is necessary since it is in practice impossible for human beings (individually or collectively) to deal with a large body of human beliefs such as the scientific corpus in the massively open-ended manner that an ideal Bayesian subject would be capable of.

In summary, neither of the two major alternatives to the approach introduced in Sect. 2 seems to be tenable.

4 Values in applied science

As clarified above, with "applied science" I mean here science that is used to guide practical decision-making. Practical uses can influence both what data we acquire and how we interpret the data that is available to us. In line with the delimitations stated in Sect. 1, I will deal here only with how practical uses that we can foresee influence our interpretations of data.

It is essential to distinguish, in all such contexts, between the practical decision to be guided by science and the scientific process as such. Often two decisions have to be based on the same scientific information: the intrascientific decision concerning what to believe and an extrascientific (practical) decision concerning what to do. As an example of this, consider a science-based decision whether or not to restrict the use of a certain preservative agent in baby food due to indications that it may disturb the production of growth hormone. The available scientific evidence will have to be assessed to determine whether a restriction is warranted. The same information will also be assessed by scientists in order to determine whether or not the presence of this effect should be considered as scientifically established. These are two different decisions, although they make use of the same scientific data.

The discussion of values in science that began with Rudner's (1953) article has been strongly influenced by a more or less explicit ideal that (pure) science should be as independent as possible of moral values. In practical decisions based on scientific information, such a requirement would not make sense. Our practical decisions do not have to be value-free, or independent of (non-epistemic) values. This applies to those decisions for which we use scientific guidance just as it applies to those that are not based on scientific information. There is, for instance, no credible reason why our decisions on what substances we should allow in baby food should be uninfluenced by the values that determine how we protect young children against other dangers in society.

We can see from this that the value problem of pure science—namely how to avoid the influence of non-epistemic values—is not present in practical applications of science. However, another problem arises instead: science is essentially arranged according to the criteria of pure science; in particular non-epistemic values have been excluded as far as possible. In practical decisions we have to combine this value-deprived scientific information in an optimal way with the values on which we wish to base our decisions. Thus, the problem is not how to exclude extrascientific values but rather how to *include* them in the intended way. This may be called a *reversed value problem*. It is not necessarily less intricate than the more commonly discussed problem of how to make science as value-free as possible.

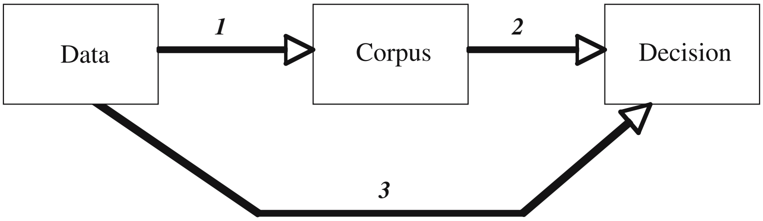

Fig. 2 The use of scientific data for decision-making (Hansson 2004)

Figure 2 illustrates the practical use of scientific information (Hansson 2004). The obvious way to use science for decision-guiding purposes is to employ information from the corpus (arrow 2). In many cases, this is all that we need to do. The high entrance requirements of the corpus have the important effect that the information contained in the corpus is reliable enough to be relied on in almost all practical contexts. Only on very rare occasions do we need, for some practical purpose, to apply stricter standards of evidence than those that regulate corpus inclusion.

However, the high entry requirements for the corpus also have another, more complicating implication. On some occasions, evidence that was not strong enough for corpus entry may nevertheless be strong enough to have legitimate influence in some practical matters. To exemplify this, consider again the case of a preservative agent in baby food that is suspected to have a negative health effect. We may assume that the evidence weighs somewhat in the direction of there being an effect, so that most scientists consider it to be more probable that the effect exists than that it does not. Nevertheless, the evidence is not conclusive, and the issue is still open from a scientific point of view. Considering what is at stake, it would be perfectly rational for a food company or a government agency to decide in such a situation to cease the use of the substance. Such a decision would have to be informed by scientific information that did not satisfy the criteria for corpus entry. More generally speaking, it would not seem rational—let alone morally defensible—for a decision-maker to ignore all preliminary indications of a possible danger that do not amount to full scientific proof. We typically wish to protect ourselves against suspected health hazards even if the evidence is much weaker than what is required for scientific proof. Therefore, in order to guide the type of decisions that we want to make, these decisions have to be based on standards of evidence that differ from the criteria used for intrascientific purposes.

Evidence that is weaker than the requirements for corpus entry cannot influence decisions in the "standard" way that is represented in Fig. 2 by arrows 1 and 2. In cases like this, we need to take a direct way from data to practical decision-making (arrow 3).

Just like the process represented by arrow 1, the bypass route represented by arrow 3 involves an evaluation of data against criteria of evidence. However, the two evaluation processes differ in being calibrated to different criteria for the required strength of evidence. The process of arrow 1 is calibrated to the standard scientific requirements, whereas that of arrow 3 is calibrated to criteria corresponding to the needs of a practical decision.

It is essential, however, that this difference does not lead, in the latter process, to inefficient use of the available scientific information. In order to see what this requires it is instructive to compare the processes represented by arrows 1 and 3. First of all, there should be no difference in the type of evidence that is taken into account. Hence, in the baby food example, the same experimental and epidemiological studies are relevant for the intrascientific decision (arrow 1) and for the practical one (arrow 3). The evidence is the same, although it is used differently. Furthermore, the assessment of how strong the evidence is should be the same in the two processes. What differs is the *required* level of evidence for the respective purposes.

In summary, when the corpus has to be supplemented (arrow 3), the same type of scientific evidence should be used, and the same relative weights should be assigned to different pieces of evidence, as in the intrascientific process of corpus formation (arrow 1). In this way the impact of non-epistemic values will be limited, as far as achievable, to the choice of an appropriate level of evidence for practical action.[6] This is proposed as part of the solution of the reversed value problem. The proposal is constructed to ensure the full use of science in the issues that science are best equipped to deal with.

In practice, decision-guiding science does not in general comply with this model. I propose, as an hypothesis, that scientific agreement is easier to achieve when as much as possible is done to isolate the scientific issues from other considerations. So-called "scientific controversies" in applied science seem to be particularly difficult to resolve when scientists are required to deal with unclear mixtures of science issues and policy issues.

In summary I propose that for practical as well as scientific purposes we should primarily use the scientific general-purpose corpus. Its high entry requirements guarantee that the information in the corpus is reliable enough to be trusted in almost all practical contexts. But from another point of view, the corpus is insufficient for many practical decisions: Due to the same high entry requirements, information that has not been included in the corpus may nevertheless have sufficient evidential weight to legitimately influence some practical decisions. In such cases, a separate evaluation of these data should be made, for the purpose of the particular practical decision. In this evaluation, the same types of evidence should be used as in intrascientific deliberations, and degrees of evidence should be assigned in the same way. The crucial difference is that the required level of evidence is different. The influence of non-epistemic values should be limited to the determination of required levels of evidence.

5 Alternative views on applied science

There are several alternatives to the approach to applied science that was presented in the previous section. An obvious alternative is to use only such scientific information for decision guidance that has been admitted into the corpus, i.e. reject the bypass route represented by arrow 3 in Fig. 2. Another alternative would be to adjust the corpus so that it suits the practical purposes, and then use this adjusted corpus as the single source of scientific information (and thus reject arrow 3 in this case as well). A third alternative would be to refrain from using the corpus at all in applied decisions and instead base them directly on the scientific data.

[6] It is important to note that this account concerns interpretations of the available evidence, and does not cover decisions on what new evidence to collect. Clearly, the chosen level of evidence can legitimately influence decisions on what types of investigations should be performed.

5.1 "Sound science"

The demand that only well-established scientific fact should be used in applied science has particularly outspoken proponents in the field of chemical risk assessment (Durodié 2003). The debate on standards of evidence in that area has often focused on the precautionary principle. This is a principle that allows regulatory decisions to be based on preliminary scientific indications that an exposure can have adverse effects on human health or the environment (Sandin 1999; Sandin et al 2002; Hansson 1999). Critics have claimed that the principle is "unscientific" and "marginalises the role of science" (Gray and Bewers 1996). The alternative view is that measures to protect the environment should only be taken if they can be based on fully established scientific fact. This has often been called the application of "sound science".[7]

The central claim of "sound science" proponents is that we should use the intrascientific standards of evidence not only for intrascientific decisions but also for the practical decisions that are based on science. However, in spite of its backing from influential organizations, this view about the relationship between the two decision processes is obviously untenable. Practical rationality demands that the required weight of evidence in practical matters be influenced by our evaluations of the possible outcomes. When there are strong but contestable indications that a volcano may erupt in the next few days, we evacuate its surroundings, rather than waiting for full scientific evidence that the eruption will take place. A principle that we should act against other types of possible dangers only if we have full scientific proof would be difficult to defend.

It should also be observed that nobody seems to apply "sound science" to all policy areas. Those who apply it to environmental measures are not necessarily prepared to wait for full evidence in all other political issues. Hence, political leaders who favour "sound science" in environmental issues tend to honour other principles in issues of national security (Consider for instance the evidence based on which the American president decided to act, in the year 2003, as if Iraq had weapons of mass destruction.).

In summary, "sound science" conflicts with practical rationality, and does not deserve to be taken as a serious alternative.

5.2 An applications-adjusted corpus

It may at first sight seem to be a good idea to adjust the corpus so that its criteria of evidence suit the practical purposes for which we intend to use it. We could for instance adjust the requirements of evidence in toxicology so that they coincide with the requirements for the practical decisions that we intend to make.[8] After such an

---

[7] This use of the term "sound science" seems to have begun with the formation of The Advancement of Sound Science Coalition (TASSC). This organization was set up by Phillip Morris company in 1993. Its major task was to promote the view that the evidence for health risks from passive smoking was insufficient for regulatory action. In spite of its name, the organization promoted pseudoscience rather than science (Mooney 2005).

[8] I am not aware of any full-blown proposal of this nature, but several authors have proposed that statistical practices should be adjusted through the use of a higher value of α^* higher than .05. Hence, Dobbins (1987: 43) proposed that it would be useful to "increase the chance of an alpha error in an individual study from the usual 5% level to 10% or greater" (cf. Cranor 1990). However, an increased value of α^* would increase the probability of type I errors (false positives), which is not in itself an advantage. A high α^* does not consistently protect us against type II errors (false negatives), which is obviously what these critics want to achieve. Hence, if an exposure kills 1 out of 1,000 exposed

adjustment has been performed, things will, presumably, be much simpler since we no longer need to distinguish between criteria for practical and intrascientific decisions.

Unfortunately, this proposal does not work for the simple reason that practical purposes differ so that there are no stable and general "practical" standards of evidence. We can assess the toxicity of a drug in relation to its use against a minor disease or against a life-threatening condition. Similarly, we may ask questions about the safety of a vaccine when it is considered for use under normal conditions or under the conditions of an extreme emergency. A question about the strength of a material can be asked by an airplane constructor or by someone designing a decoration. In general, due to the variability of our practical uses of science, there is no single, well-determined way to adjust the standards of evidence to practical purposes. A unique corpus cannot be based on the diverse criteria that we employ in practical decisions.

5.3 Doing without the corpus

The alternative not to use the corpus at all for practical decision-making does not seem to have any outspoken adherents, but it should nevertheless be treated for the sake of completeness. The idea would be to reason anew from the data in each particular case when a science-based decision has to be made. This however, would be at least as cognitively unreasonable as the Bayesian approach discussed in Sect. 3.2. It would drive us to reappraisals of an immense mass of empirical conclusions, hypotheses, and theories. Such a reappraisal could be performed by a hypothetical, ideal Bayesian subject, but it is far beyond the reach of human scientists of flesh and blood. Each of us has access only to small parts of the entire corpus of knowledge on which modern science is based. To reason anew without using the corpus would in practice be an impossible feat.

6 Conclusion

The gist of this article is a proposal that we should use a general-purpose corpus for both scientific and practical deliberations, but supplement it with additional scientific information when this is justified by the standards of evidence that we use for particular, practical purposes. The major principles for the corpus and its supplementations can be summarized as follows:

1. *We need a general-purpose corpus of science*. As decision-makers and cognitive agents with limited capacity, we need a general-purpose, intersubjective, and continually updated corpus that summarizes the available scientific information.
2. *The information in the corpus should be sufficiently reliable to be used for all or nearly all purposes.* This is in practice best achieved by using the strict standards of evidence in intra-scientific deliberations as entry requirements for the corpus.
3. *The corpus should not be adjusted to particular practical purposes or values.* A general-purpose corpus should not preferentially serve certain values. This too

Footnote 8 continued
persons, then a study on 200 persons of whom no one was killed does not give us the information we need. This cannot be remedied by any choice of α^*. What needs to be done, instead, is to keep track of the statistical power of the tests. Statistical power is closely related to the risk of false negatives (Hansson 1995).

is achieved by adjusting the corpus to scientific criteria of evidence, since these programmatically exclude considerations of practical value.

4. *The corpus should be supplemented with additional scientific information whenever this is justified by the evidential requirements of a practical problem.* Since we have adopted high entry requirements for the corpus, information that is not included in the corpus may nevertheless be sufficiently reliable for some practical purposes. As a major example of this, indications of a danger may be sufficiently strong to justify action against the danger even if they do not satisfy the requirements for inclusion into the corpus.
5. *In the assessment of such supplements to the corpus, the strength of the available evidence should be judged according to the same scientific criteria that are used in the formation of the corpus.* Hence, considerations of practical value should be excluded, as far as possible, from the determination of how strong the evidence is. Instead, the role of practical values is to determine how strong the evidence must be for the different practical decisions that are at stake.

References

Cranor, C. F. (1990). Some moral issues in risk assessment. *Ethics, 101*, 123–143.
Dobbins, J. G. (1987). Regulation and the use of 'negative' results from human reproductive studies: The case of ethylene dibromide. *American Journal of Industrial Medicine, 12*, 33–45.
Durodié, B. (2003). The true cost of precautionary chemicals regulation. *Risk Analysis, 23*(2), 389–398.
Feleppa, R. (1981). Epistemic utility and theory acceptance: Comments on Hempel. *Synthese, 46*, 413–420.
Gray, J. S., & Bewers, M. (1996). Towards a scientific definition of the precautionary principle. *Marine Pollution Bulletin, 32*(11), 768–771.
Hansson, S. O. (1995). The detection level. *Regulatory Toxicology and Pharmacology, 22*, 103–109.
Hansson, S. O. (1996). What is philosophy of risk? *Theoria, 62*, 169–186.
Hansson, S. O. (1999). Adjusting scientific practices to the precautionary principle. *Human and Ecological Risk Assessment, 5*, 909–921.
Hansson, S. O. (2004). Philosophical Perspectives on Risk. *Techne, 8*(1), 10–35.
Hansson, S. O. (2006a). Falsificationism falsified. *Foundations of Science, 11*, 275–286.
Hansson, S. O. (2006b). Praxis relevance in science. *Foundations of Science* (in press).
Harsanyi, J. C. (1983). Bayesian decision theory, subjective and objective probabilities, and acceptance of empirical hypotheses. *Synthese, 57*, 341–365.
Hempel, C. G. (1960). Inductive inconsistencies. *Synthese, 12*, 439–469.
Jeffrey, R. C. (1956). Valuation and acceptance of scientific hypotheses. *Philosophy of Science, 23*, 237–249.
Levi, I. (1962). On the seriousness of mistakes. *Philosophy of Science, 29*, 47–65.
McLaughlin, A. (1970). Science, reason and value. *Theory and Decision, 1*, 121–137.
Mooney, C. (2005). *The republican war on science*. New York: Basic Books.
Peirce, C. (1934). The fixation of belief. In C. Hartshorne, & P. Weiss, (Eds.), *Collected papers of Charles Sanders Peirce*, Vol. 5 (pp. 223–247). Harvard: Harvard University Press.
Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science, 20*, 1–6.
Sandin, P. (1999). Dimensions of the Precautionary Principle. *Human and Ecological Risk Assessment, 5*(5), 889–907.
Sandin, P., Peterson, M., Hansson, S. O., Rudén, C., & Juthe, A. (2002). Five charges against the precautionary principle. *Journal of Risk Research, 5*, 287–299.
Ziman, J. (1996). 'Postacademic science': Constructing knowledge with networks and norms, *Science Studies, 9*, 67–80.

Biography

Sven Ove Hansson is professor of Philosophy at the Royal Institute of Technology, Stockholm. He is also a member of the Swedish government's advisory board of researchers. His major research areas are philosophy of risk, decision theory, epistemology, belief dynamics, and value theory. His most recent books are *Setting the Limit. Occupational Health Standards and the Limits of Science* (Oxford UP 1998), *A Textbook of Belief Dynamics. Theory Change and Database Updating* (Kluwer 1999), and *The Structures of Values and Norms* (Cambridge UP 2001). Homepage: http://www.infra.kth.se/~soh.