

# Remarks on Hansson’s model of value-dependent scientific corpus

Philippe Stamenkovic\*

3rd March 2023

## Abstract

This article discusses Sven Ove Hansson’s corpus model for the influence of values (in particular, non-epistemic ones) in the hypothesis acceptance/rejection phase of scientific inquiry. This corpus model is based on Hansson’s concepts of scientific corpus and science ‘in the large sense’. I first present Hansson’s corpus model of value influence with some introductory comments about its origins, a detailed presentation of the model with a new terminology, an analysis of its limits, and an appreciation of its handling of controversial non-epistemic values. I conclude that it is a very good candidate for managing value influence, because contrary to other models in the literature it is fairly simple, it provides a universal and systematic procedure for dealing with values, and it systematically preserves the epistemic integrity of science. I then study some difficulties associated with the model’s central feature of taking the maximum level of evidence required by the applications of a claim in order for the latter to enter the scientific corpus: the difficulty to identify this maximum requirement; the non-optimality of this requirement with respect to other less demanding requirements; the potentially detrimental preference for false negatives over false positives on which it relies; and the issue of whether there can nevertheless be requirements higher than this maximum. I show that these issues may potentially challenge the corpus model. Finally, I call for empirical work (in particular, an investigation of scientists’ and engineers’ own normative views) in order to better assess these issues.

**Keywords:** corpus model, Sven Ove Hansson, non-epistemic values, level of evidence, scientific corpus.

**Mots-clés:** modèle de corpus, Sven Ove Hansson, valeurs non épistémiques, niveau de preuve, corpus scientifique.

## 1 Introduction

In the philosophy of science, ‘non-epistemic’ (also called ‘non-cognitive’, or ‘contextual’) values designate all those extra-scientific values (such as social, moral, economic or political values) which are not concerned with furthering the truth-seeking goal of science<sup>1</sup> and do not intrinsically

---

\*philippe.stamenkovic@icloud.com.

<sup>1</sup>Note that truth can be considered either as (one of) the defining aim(s) of science, if considered in itself; and/or as a value, if considered in relation to us, who value it (Hicks, 2014, pp. 3272–3273). Hempel (1981, p. 404) goes as far as to consider that epistemic values can also be considered as constitutive aims of science, instead of ‘conceptually independent means for its attainment’.

belong to scientific activity, in contrast to ‘epistemic’ (or ‘cognitive’, or ‘constitutive’) values, which designate the intra-scientific values which contribute to this goal (such as empirical adequacy, external and internal coherence, simplicity, unity, fruitfulness, etc.) (e.g. Kuhn, 1977; McMullin, 1982). While non-epistemic values in science can threaten the epistemic autonomy and integrity of science, they can also contribute to a more socially responsible science. Although considerations about epistemic values in science are practically as old as science itself, a proper debate about values (especially non-epistemic ones) in science really begins after World War II (Churchman, 1948; Rudner, 1953; Jeffrey, 1956; Levi, 1960; Hempel, 1960). Until the end of the 20th century, the value-free ideal (VFI) of science, the normative position according to which there *should* be no influence of non-epistemic values<sup>2</sup> (even if there *is* such in practice), is the dominant view among philosophers and especially scientists. Hempel (1960; 1981) is an early representative of this view (as well as of the corpus model of scientific knowledge), according to which only ‘*purely scientific, or epistemic, utilities*’ (1960, p. 465) should influence our decision to accept or not a hypothesis (see section 2.1). In the last decades however, the VFI has been the object of sustained critique (e.g. Longino, 1990; Douglas, 2009; Kourany, 2010; K. Elliott, 2017), to the point that its rejection – which we may dub the value-laden ideal (VLI) – now appears as the consensus view in the philosophy of science. For example, Holman and Wilholt (2022) consider the VFI/VLI debate over, and consider the new debate (which they dub ‘the new demarcation problem’) to be about, not whether or not to allow influence of non-epistemic values, but how to distinguish legitimate from illegitimate forms of such influence.

This debate reflects on the appropriate influence of values in the various phases of scientific inquiry, concerning: 1) what to investigate (definition of research avenues); 2) how to investigate it (gathering of evidence and choice of methods, including from the ethical point of view); 3) what to conclude from our investigations (acceptance or rejection of a hypothesis or theory); 4) how to use and communicate the results of our investigations<sup>3</sup>. Another way of analysing the debate is to list the following arguments in favour of values in science (Elliott, 2011, p. 62; 2022, §3), which roughly correspond to phases 2 and 3 above (except the first one which is focused on phase 3 only):

- the inductive risk (or ‘error’) argument: in accepting or rejecting a hypothesis, a scientist has to consider the risk of being in error, by either wrongly accepting an actually false hypothesis (‘false positive’) or wrongly not accepting an actually true hypothesis (false negative)<sup>4</sup> (e.g. Douglas, 2009);
- the underdetermination (or ‘gap’) argument: inherently value-laden concepts and background knowledge are used by scientists to connect theory and evidence (e.g. Longino, 1990)<sup>5</sup>;

---

<sup>2</sup>This is the restricted sense of the VFI, concerning only the acceptance/rejection phase. In the large sense of the VFI, there should be no influence of non-epistemic values on (some of) the other phases as well (except regarding morally acceptable ways of performing research), typically phase 1. This large acceptance of the VFI has few supporters in the philosophy of science, although it can have popularity in practising scientists.

<sup>3</sup>Some authors talk of the contexts of discovery, justification (following the widely held distinction) and application to refer to phases 1 and 2, 3, and 4 respectively (e.g. Bueter, 2021); or of pre-epistemic (1 and 2), epistemic (3), and post-epistemic (4) phases (Hicks, 2014).

<sup>4</sup>In statistical theory, false positives are called type I errors and false negatives type II errors. Their rigorous definition is based on the null hypothesis, according to which the two possibilities investigated are the same (any difference is due to chance only). Type I error amounts to wrongly rejecting an actually true null hypothesis, and type II of wrongly not rejecting an actually false null hypothesis.

<sup>5</sup>According to critiques (outside the scope of this article), the error argument can be seen as a special case of the gap argument (e.g. ChoGlueck, 2018). Here I will leave aside the issue of the error/gap distinction. Whatever the exact stance taken with respect to this distinction, it is clear that the gap argument is a larger, more abstract issue

- the aims argument: in order to achieve the non-epistemic aims of science of obtaining knowledge which is relevant for its users, scientists have to take into account non-epistemic values to assess the adequacy of their models, hypotheses and theories;
- the conceptual argument: non-epistemic values are relevant to assessing hypotheses that include concepts which mix epistemic and non-epistemic content.

All these issues belong to the ‘new demarcation problem’<sup>6</sup>. It is especially the influence of non-epistemic values on phase 3, where hypotheses and theories are accepted or rejected, which is controversial. For the other phases (1, 4 and to a lesser extent 2), the recourse to non-epistemic values is generally taken to be uncontroversial. This acceptance/rejection phase is the only one which will be considered here. In the following, the term values without further specification designates all kinds of values (epistemic as well as non-epistemic, even if the major concern is about the latter), and the expression influence of values without further specification refers to their influence on the acceptance/rejection phase.

The expression VLI includes in fact a variety of positions allowing different types of non-epistemic value influence (for useful classifications of different types of value influence, see Holman and Wilholt, 2022; Kevin C. Elliott, 2022)<sup>7</sup>. But whatever the position within the VLI, it is confronted to difficulties (see e.g. Reiss and Sprenger, 2020; Kevin C. Elliott, 2022), and while these difficulties have not been settled, the question of values in science remains controversial in the philosophical literature (K. Elliott and Steel, 2017). While allowing non-epistemic value influence seems not only inevitable but also desirable, a consensual model for doing so remains to be seen.

It should be noted that the issue of non-epistemic value influence exceeds the narrow confines of philosophy of science, and concerns several aspects of the science-society relationship. Apart from its obvious relevance for policy-making, it gives rise to several problems of public trust in science, due for example to suspicions of biased science or fears of technocratic value judgements passed by experts (Carrier, 2022). It is reasonable to assume that the controversial and complicated nature of the debate taking place in the philosophy of science, which has become largely out of reach not only for the general public and policy-makers, but also for scientists not specialising in this sub-field of philosophy of science, does not help remedy this situation. Rather, this situation allows a confused image of science to arise in the view of policy-makers, the wider public, and even scientists – with a potential negative impact on public trust in science (for a study of the contrasted effects of values in science on the general public’s trust, see Kevin C Elliott et al., 2017). Worse, such unsettled debate can be used by science deniers to support their undermining enterprise (Hansson, 2020a, p. 37). For example, non-epistemic values have been, and can be used to put into question policy-relevant scientific claims with big social stakes, such as typically anthropogenic global warming or the negative environmental and/or health effects of pesticides. Therefore, progress on this issue is urgently needed.

The philosophical debate generally provides increasingly sophisticated guidelines for acceptable (or even beneficial) influence of non-epistemic values only on a *case-by-case* basis, and/or by *combining* different approaches (e.g. Kevin C. Elliott, 2022). Like many other overspecialised but endlessly controversial areas of philosophy, the debate may sometimes look like a ‘self-indulgence

---

than the error argument. Here I will focus on the latter, which is the object of Hansson’s model. For a critique of the underdetermination argument, see Ruphy (2006), an underestimated proposal, like Hansson’s.

<sup>6</sup>Another strand of literature (e.g. Anderson, 2004; Ruphy, 2006) also considers the reverse influence of evidence on values, in other words how value judgments themselves can be empirically tested. Hansson does not consider this possibility, and takes values as inputs in his model.

<sup>7</sup>Absent from these classifications, however, is Hansson’s corpus model (see below).

for the few’ (Kitcher, 2011, p. 248), where no progress is really made and no consensus is ever achieved. On the contrary, it seems that in order to exit our current predicament we need a:

1. reasonably *simple*,
2. *universal* (i.e. applicable to all cases, in contradistinction to case-by-case models) and
3. *systematic* (i.e. taking into account all possible cases in a pre-determined way) model of how values should be allowed in science,
4. while at the same time systematically guaranteeing the epistemic *integrity* of science (more about this concept in section 2.2.1), which is the main concern when allowing non-epistemic values in science, and the chief argument used by VFI defenders<sup>8</sup>.

But in fact, there is such a candidate model for non-epistemic value influence, at least with respect to the inductive risk issue<sup>9</sup>, which is regrettably neglected in the debate about values in science. Indeed, Hansson (2007; 2010; 2014; 2017; 2018; 2020) has progressively developed and refined a model of (epistemic as well as non-epistemic) value influence which is, to my knowledge, the only one to be (reasonably) simple, universal, systematic, and, most importantly, which systematically preserves the integrity as well as the unity of science when non-epistemic values are allowed in. The purpose of this article is to present in detail this ‘corpus model’ of value influence (so called by me because it is based on Hansson’s concept of scientific corpus), identify its strengths and weaknesses, and review potential difficulties associated with it. Section 2 introduces Hansson’s corpus model with some preliminary comments on its origins (§2.1), presents the model in detail (§2.2), analyses its limits (§2.3), studies how well it handles controversial non-epistemic values (§2.4), and summarises its advantages (§2.5). Section 3 reviews some potential difficulties associated with the model’s central feature of taking the maximum level of evidence required by the applications of a claim in order for the latter to enter the scientific corpus: difficulty to identify this maximum requirement (§3.1); non-optimality of this requirement with respect to other less demanding requirements (§3.2); the potentially detrimental preference for false negatives over false positives on which it relies (§3.3); and the issue of whether there can nevertheless be requirements higher than this maximum (§3.4). Section 4 concludes.

## 2 Presentation of the corpus model

To account for (epistemic as well as non-epistemic) value influence (the two types of values are treated indistinctly by the model), Hansson has developed and refined what I will call his ‘corpus model’<sup>10</sup>, in the course of several publications (2007; 2010; 2014; 2017; 2018; 2020). In a nutshell, the idea is, on the basis of a universal, multi-purpose scientific corpus, to require a unique level

---

<sup>8</sup>I agree that a complicated, case-by-case model could also ensure the epistemic integrity of science, and reach a consensus in the philosophy of science. The main problem, however, would be its perception outside philosophy of science, especially in terms of public trust in science. It is reasonable to assume that a simple and universal model would be beneficial in this respect. For a discussion of simplicity, see footnote (41).

<sup>9</sup>To my knowledge, Hansson does not address the wider underdetermination issue, but instead focuses on the inductive risk argument, for which his corpus model (see hereafter), based on the level of evidence required to accept a claim, is particularly adapted.

<sup>10</sup>Hansson sometimes uses the expression of ‘corpus model’ (e.g. in 2018, *passim*) to designate his conception of the scientific corpus in general, including, but not limited to, the way new claims are incorporated in it. Here I use this expression specifically for this latter aspect, and I will talk of the corpus *concept* to designate the former.

of evidence for the acceptance of a claim into the corpus, namely the highest level required by all possible applications of this claim. It is important to note that the corpus model should be considered more or less *normative*, although Hansson’s position on this matter is not completely clear. On the one hand, Hansson hopes to provide not only a desirable, but also a realistic account of science (and indeed he has successively modified that account in order to be more realistic, as we will see in section 2.1); he is perfectly aware that conceptualising has to take into account empirical adequacy with scientific practice to some extent. For example, the assumption of a unique scientific corpus instead of several, discipline-based corpora<sup>11</sup>, is based on the observation of the increasing interdisciplinary collaborations between academic disciplines. On the other hand, Hansson does not pretend that his account correctly or fully describes actual scientific practice, and he underlines its normative nature (2018). This issue is complicated by the fact that scientific practice itself obeys to norms: the very requirement of setting a level of evidence, which is the central aspect here, cannot escape being normative to some extent. Hansson’s model can thus be read as a new proposal for this norm. Perhaps a good way to characterise it is to present it as an *idealisation* of scientific practice, based on some crucial simplifications (such as an all-purpose scientific corpus instead of multiple corpora for each application, see section 3.2), half-way between a purely descriptive account and a normative conception. In this section I first make some introductory comments on the origins of the corpus model (§2.1), before presenting it in detail with the help of a new terminology (§2.2), analysing its limits (§2.3), studying how well it handles controversial non-epistemic values (§2.4), and showing its advantages (§2.5).

## 2.1 Origins of the model

### 2.1.1 Precursor views

Hansson seems to have taken inspiration for his corpus model of scientific knowledge from Hempel (1960; 1981), whom he regularly mentions. To my knowledge Hempel is the first to mention the scientific ‘*corpus*’ as the ‘body of scientific knowledge at a given time’, which is ‘represented by the set of all statements accepted by science at that time’ (1960, pp. 462–464). Accepting a hypothesis thus means incorporating it in the corpus, although corpus membership is always provisional in principle, subject to possible subsequent discovery of disconfirming evidence which would lead to the exclusion of this hypothesis from the corpus. As seen in the introduction, Hempel tolerates only ‘*purely scientific, or epistemic, utilities*’ to have an influence on the decision to accept or not a hypothesis: such 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’ (1960, p. 465). Here Hempel uses a quantitative concept of utility<sup>12</sup> (not to be confused with the qualitative concept of value used in this article and in the literature on values in science, which designates qualities such as simplicity, empirical adequacy, etc.) which is a function assigning a real number to a possible outcome of a given action (here, accepting or not a given hypothesis), in the scope of Carnap’s rational decision-making rule of maximising the estimated utility (see Hempel, 1960, §6 for details). Hempel initially proposes a utility function for adding a hypothesis  $h$  to the corpus  $e$  ‘directly proportional to the amount of new information provided

---

<sup>11</sup>Note that this possibility is different from multiple, *application*-adjusted corpora. The latter is also excluded by Hansson, but for reasons which seem more normative than the former.

<sup>12</sup>Later he also speaks of ‘epistemic values’ (1981, p. 398), although it is unclear in the only sentence where they are mentioned whether epistemic values are strictly equivalent to epistemic utilities.

by  $h$ , or to the negative value of that amount<sup>13</sup>, according as  $h$  is true or false; and inversely proportional to the amount of information already contained in  $e'$  (466) – in other words, it only reflects the truth (or falsity) *content* of the hypothesis. He proposes a hypothesis acceptance rule accordingly (466), only to realise that it is too ‘lenient’ since it does not take into account other important epistemic values of scientific hypotheses, such as their explanatory or predictive power, their empirical adequacy, the gain in simplicity which they enable for the corpus, their scope of application, compatibility with neighbouring theories, all values which yet have to be ‘given clear and precise definitions’, as well as be weighted against each other (1960, p. 467; see also 1981, pp. 399–401). While in 1960 Hempel still seems optimistic about the precise formulation of epistemic utility, in 1981 he doubts its feasibility. As far as I know, the precise quantification of such values has never been successfully proposed, and it seems doubtful that a formal quantitative rule could ever be proposed for scientific hypothesis acceptance<sup>14</sup> – even the quantification of purely content-based utilities seems out of reach (how could the information contained in a scientific hypothesis be quantitatively evaluated?).

### 2.1.2 Hansson’s prior view

It is interesting to note that Hansson’s position has switched from the VFI to the VLI. In 2007a he still holds the traditional view according to which the scientific corpus should be ‘programmatically’ free from non-epistemic values (even if he already conceives it as a multipurpose and unique corpus), and that non-epistemic values should only influence the level of evidence needed for practical decisions. This position seems to be linked to his distinction between *pure* science – where he apparently locates the *production* of the scientific corpus, and where non-epistemic values are not allowed to influence the corpus – and *applied* science – where he apparently locates the *use* of the corpus for practical decisions, and where non-epistemic values are allowed to influence this use (not the corpus)<sup>15</sup>. This localisation of knowledge production within pure science only, which is ‘science aiming at knowledge per se’ (2007, p. 257), stands somewhat in tension with the idea of a multipurpose corpus, apt to any use (including practical use). It may come from Hansson’s lumping together *applied science* and *application of science*. Indeed he defines applied science alternatively:

- as ‘science that aims at [...] knowledge for practical use’ – which is applied science proper;
- and as ‘practical uses [or ‘practical applications’] of science’, or ‘science applied to practical decisions’ – which designates in fact applications of science (2007, pp. 257, 262).

There is indeed a difference between applied science (such as applied mathematics, engineering science, medical science) and application of science: ‘The former is a part of knowledge *production* [even with a practical aim in mind], the latter is concerned with the *use* of scientific knowledge and

---

<sup>13</sup>Thus Hempel attributes the same (positive or negative) ‘utility’ or ‘value’ (in Hempel’s quantitative sense) to a true positive (a rightfully accepted claim) and a false positive (a wrongfully accepted claim).

<sup>14</sup>As Hempel (1981, p. 402) remarks, there are some acceptance rules for specific hypotheses which can be formulated precisely and are (more or less) acknowledged in the scientific community (such as for measurements or statistical hypotheses, even if in the latter case the critical levels are largely conventional and can be discussed). But for ‘more comprehensive’ hypotheses (not to speak of theory changes), there are no generally accepted rules.

<sup>15</sup>In spite of the fact that Hansson (2007a, p. 258) does not strictly separate pure from applied science, but rather underlines the ‘problematic nature’ of this division, and recommends 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’.

methods for the solving of practical problems of action (e.g., in engineering or business), where a scientist may play the role of a consult[ant]’ (Niiniluoto, 1993, 9, my italics)<sup>16</sup>. Lumping together applied science (where knowledge is produced) and applications of science (where it is not) may lead to exclude any knowledge production from applied science proper<sup>17</sup>.

### 2.1.3 Hansson’s current view

However, from 2014 Hansson acknowledges and defends the view that non-epistemic values can and should have an influence on the scientific corpus – maybe because he concomitantly acknowledges that corpus production also takes place in applied science and that the corpus incorporates applied scientific knowledge. He apparently still lumps together applied science and application of science, but now since corpus production also belongs to applied science/applications of science, he has to take into account applications (i.e. values) in the corpus production. Hansson (2014) considers an increasing intertwinement of pure and applied science and associated corpus production (which takes place under their combined influence), and nicely illustrates how this intertwinement is paralleled by an incorporation of non-epistemic values.

Hansson’s corpus model extends and complements Hempel’s conception in several respects. The corpus model is based on Hansson’s concepts of scientific corpus and ‘science in the large sense’, which he has developed concomitantly. The scientific corpus is basically the total body of (accepted, consensual) scientific knowledge, in the form of the published<sup>18</sup> scientific literature (articles and textbooks) (see e.g. Hansson, 2018a, pp. 68–71). The corpus is interdisciplinary, all scientific disciplines collaborating with, and respecting each other; it is universal hence there is only one corpus; and it is apt to any (theoretical or practical) application since it represents our most reliable (although always provisory) knowledge (e.g. Hansson, 2007a). This conception of the scientific corpus is intimately linked to Hansson’s conception of ‘science in the large sense’, which includes all the academic disciplines in pure and applied science: the formal, natural, social, human and technological sciences (e.g. Hansson, 2018a, pp. 60–65).

I believe that Hansson’s concept of universal, multi-purpose corpus is the key to his model of value influence. Once this universality and multi-purposiveness is acknowledged, it becomes necessary to consider the theoretical and practical consequences of our choice of incorporating a claim into the corpus, since the latter can be used for any purpose, whether theoretical (e.g.

---

<sup>16</sup>We can add two precisions:

– there can be applications of pure science, not only of applied science: typically many parts of logic, or pure mathematics (such as arithmetics, analysis, geometry, probability) are used directly in almost any branch of science, as well as in many everyday applications.

– there can be theoretical (including epistemic) applications of science, not just practical ones: in addition to epistemic applications (such as in the previous paragraph), scientific knowledge can also be used for the solving of theoretical (not just practical) problems of action (e.g. which research avenues to pursue), which belong to theoretical decision-making.

<sup>17</sup>What may have prevented Hansson (2007a) from incorporating non-epistemic values is that, facing the impossibility of uniquely adjusting the corpus to its many possible applications (266), he hadn’t come up yet with his solution of taking the most demanding application for setting the final level of evidence required for corpus entry, a solution which first appears in Hansson (2014) and has the advantage of covering all possible applications while ensuring the epistemic integrity of science at the same time, as we shall see.

<sup>18</sup>Note that in practice, unpublished scientific literature (even if it is good science) is not part of the corpus by definition, since it cannot be (widely) used by others than the author(s). Conversely, bad (false or fake) science may also inappropriately be part of the corpus, although it should (hopefully) later be rejected. Therefore the scientific corpus is not actually equivalent to scientific *knowledge*: there might be true claims not part of the corpus, and false claims part of it. But these considerations are put aside in the idealised model under consideration here.

for pursuing research) or practical (e.g. for manufacturing goods, making policy, or adjudicating court cases (ChoGlueck, 2018, p. 713)). Therefore, it becomes necessary to take into account (epistemic as well as non-epistemic) values. Furthermore, case-by-case or combined models, like those depending on the goals of scientific activity or combining several approaches, are ruled out<sup>19</sup>. The idea of a universal (hence unique), multi-purpose corpus itself comes from Hansson’s acknowledgment of our cognitive limits and cognitive economy needs, which is, in my opinion, a very original and convincing feature of his conception. Indeed, Hansson insists on the cognitive economy allowed by a universal social repository of factual statements, which can be used by anyone for any kind of purpose (including practical purposes associated with risk). As Hansson (2018a, 71, my translation) claims: ‘[...] the [scientific] corpus is universal, i.e. we use the same set of temporarily fixed scientific beliefs for all sorts of ends. [...] The need for universality of the scientific corpus comes from individual cognitive economy rather than pressure for social cooperation. A community of beings without cognitive limits or with insignificant cognitive limits can have no difficulty to apply different belief norms in different domains, if they have reasons to do so. For us, humans, such practice is unmanageable. The corpus must be universal in order to fulfil its function.’ It is the same reason which makes Hansson reject a Bayesian model of scientific corpus, in which the claims are ascribed degrees of acceptance in the form of probabilities. Such a corpus would be impossible to use for humans<sup>20</sup>. Instead, according to Hansson claims must be either accepted or rejected in a binary way.

Because the corpus can be used both for theoretical and practical purposes, it is the same corpus which is used in the context of research and in the context of policy-making. Hansson’s model therefore blends the roles of scientists as researchers (producing scientific knowledge in academic settings) and scientists as experts (providing users of scientific knowledge such as decision-makers with *applicable* scientific knowledge), illustrating to some extent what Gundersen (2018, p. 53) calls the ‘inseparability view’ regarding the roles of scientists as researchers and as experts<sup>21</sup>. In practice, the responsibility to contemplate all the consequences (including practical ones) of incorporating a claim in the corpus falls to scientists, since they are the ones who decide to accept a claim or reject it, and go through the procedure of the corpus model described in section 2.2. As we shall see in section 3.1, this forecasting can be quite difficult in practice.

Finally, note that Hansson’s model does not deny the possibility nor the reality of value-free science (contrary to conceptions such as Longino’s (1990)): if there are no (practical, but if relevant also theoretical<sup>22</sup>) applications of a claim, then there is no influence of non-epistemic values on this claim (even if in principle their influence is always vindicated). Indeed, there are many parts of science – and even entire disciplines or sub-disciplines –, in formal, natural or human science,

---

<sup>19</sup>Again, if we follow Hansson’s *normative* conception of a multi-purpose corpus; in practice, case-by-case models might well be used.

<sup>20</sup>The fact that we do accept hypotheses all the time (in everyday life as well as in science), and the need to supplement Bayesian decision theory by ‘a theory of tentative acceptance of empirical hypotheses’, were already acknowledged by Harsanyi (1983, p. 341).

<sup>21</sup>Gundersen’s article insists on the normative commitments of the *roles* of scientists as researchers and as experts, whereas Hansson’s conception is about the *product* of scientific research (the corpus). Nevertheless, Gundersen’s conceptual distinction in fact also applies to the latter: ‘In this [inseparability] view, scientists are equally responsible for the consequences of their *results* in the context of research and in the context of policymaking.’ (Gundersen, 2018, 53, my italics)

<sup>22</sup>One would naturally think that non-epistemic values only play a role for practical applications, but they can also play a role for some theoretical applications (such as the choice of some research directions), although not necessarily for all (hence the relevant qualification). On the other hand, it is hard to imagine practical applications not relying on some non-epistemic value(s). Furthermore, there will almost always be epistemic value influence for the acceptance of a claim.



which do not have any known or foreseeable practical application (and are indeed dedicated to pure knowledge, hence can be called pure science), such as parts of pure mathematics, physics (e.g. cosmology), or philosophy (e.g. metaphysics). In such areas non-epistemic values may be – and often are – completely irrelevant. Nevertheless, the ideal supported by Hansson is that of value-laden science. Put differently, as soon as an application is known or thought of, it must be taken into account and allowed to influence the corpus; doing differently would be theoretically or practically (e.g. morally, socially or politically) reprehensible. Thus, values can have an influence not only on applied science, but also pure science (which may always be considered with a theoretical or practical goal in mind, although this goal was not originally present).

## 2.2 Presentation of the model

Recall that we are concerned only with the acceptance/rejection (or inductive risk) phase of scientific inquiry, which is indeed what the corpus model is about. Before presenting the corpus model, I will define some key terms which will be helpful for the following. This terminology is somewhat different from the one Hansson (2014, p. 133; 2017, p. 215) uses<sup>23</sup>, but I believe it will contribute to clarify and simplify his model.

### 2.2.1 Preliminary terminology

Firstly, regarding the *decisions* that we can take, we can distinguish the following categories:

- *Theoretical* decisions concern knowledge: they are made up of two distinct subcategories:
  - *epistemic* decisions concerning our choices of what to believe, i.e. the acceptance or rejection of a claim according to the available evidence;
  - *non-epistemic* (theoretical) decisions concerning our choices of what to do in order to achieve theoretical aims (related to the pursuance of knowledge), i.e. theoretical action.
- *Practical* decisions concern our choices of what to do in order to achieve practical aims (not related to knowledge), in other words practical action. Practical decisions are all non-epistemic.

In science, only theoretical decisions are taken<sup>24</sup>. They can be either epistemic theoretical decisions (when we accept or reject a claim) or non-epistemic theoretical decisions (during all our scientific endeavours, for example when we choose research avenues, or when we decide to perform actions in order to gain information). Both types of theoretical decisions can be imbued with values (either epistemic or non-epistemic values). The category of non-epistemic decisions thus includes non-epistemic theoretical decisions, as well as practical decisions, as illustrated by figure 1.

Secondly, regarding the *levels of evidence* which we require to take such decisions with respect to a given claim, we can distinguish:

- the level of evidence required for the epistemic decision to accept this claim, in other words for its entry into the scientific corpus: let us call this the *epistemic requirement* for this claim;

---

<sup>23</sup>For example, Hansson doesn't talk of epistemic *decisions*, even if that is in fact what they are (since they concern the decision to accept or reject a claim).

<sup>24</sup>Practical decisions can be taken in the *applications* of science, see section 2.1.2.

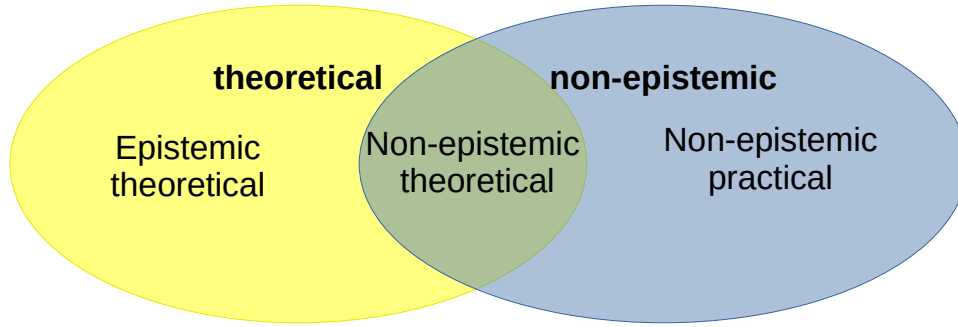


Figure 1: Different types of decisions.

- the level of evidence required to take a non-epistemic (theoretical or practical) decision on the basis of this claim: let us call this the *non-epistemic requirement* for this claim.

Note that the corpus model presupposes:

- obviously, that you can *assign* a specific level of evidence required for a claim to be accepted in the corpus;
- in addition, that for a given claim the levels of evidence required for various (theoretical or practical) decisions can be at least *sorted* (if not *quantified*), since the model takes the maximum of these levels.

Finally, let me briefly reflect on the concepts of *integrity* and *reliability*, two concepts used (somewhat indistinctly) but not defined by Hansson. I take integrity to be a purely epistemic concept which insists on the absence of undue influence (illegitimate value influence, bias and other distorting factors) in knowledge, or more precisely knowledge production. It is not exactly equivalent to reliability, which seems to be a wider concept (applicable rather to knowledge itself than to its production process), both epistemic (something related to truthfulness) and (even more so) non-epistemic, oriented towards application (whether theoretical or practical). I don't intend to analyse these concepts in detail here. Let me just note that it seems *prima facie* fair to say that breach of integrity makes unreliability more probable. Nevertheless, integrity does not necessarily imply (is not sufficient for) reliability: results of research performed in an integer way may nevertheless not be reliable (because they lack some other essential quality, such as empirical adequacy for example). Conversely, we might think of reliable research results which nevertheless were not obtained in an integer way, by a stroke of luck.

## 2.2.2 Presentation of the model

In this subsection all affirmations, claims and (normative) conceptions are those of Hansson (except for some terminology modifications and some comments in the footnotes). Again, they are just reformulated and recapitulated here. Now, as Hansson (2014, p. 132) remarks, in general 'that which we classify as knowledge can appropriately be relied upon in practical reasoning and action'. With respect to scientific knowledge in particular, the level of evidence required for a claim to enter the scientific corpus (i.e. for its acceptance) on purely epistemic grounds, which is high, is 'reliable

enough for the vast majority of [theoretical as well as] practical purposes'. We can generally use scientific knowledge and rely on it as a basis for our theoretical and practical decisions. However, it may happen that this level of evidence is, with respect to the non-epistemic theoretical or practical decisions that we may take on the basis of this claim, either:

- too high (in other words, *science demands too much proof* for our non-epistemic goals): for example, we may want to perform research based on hypothetical acceptance of a hypothesis even if it is not accepted in the scientific corpus (theoretical decision); or we may want to take action in the face of a natural or technological dangerous event even if its occurrence is not scientifically confirmed (practical decision);
- or too low (*science demands not enough proof*): for example, including a false statement in the corpus may have particularly negative theoretical impact (e.g. in terms of effects on future scientific endeavours) or negative practical impact (e.g. in terms of safety of a device or process).

In such cases of discrepancy between epistemic and non-epistemic requirements, for the claim in question we can modify:

- either its epistemic requirement, i.e. the level of evidence required for its entry into the corpus: this is what Hansson (2014, p. 133) calls an 'epistemic adjustment' of this level;
- or its non-epistemic requirement, i.e. the level of evidence required for non-epistemic (theoretical or practical) decision-making on the basis of this claim: let us call this a 'non-epistemic adjustment' of this level<sup>25</sup>. In this case the epistemic requirement remains unchanged. (Note that for the same scientific claim, non-epistemic adjustments can differ according to the envisaged application<sup>26</sup>.)

Both adjustments (epistemic and non-epistemic) can in principle be made in both directions: upward i.e. higher level of evidence required; or downward i.e. lower level of evidence required. However two types of adjustment are problematic:

- downward epistemic adjustments, because they threaten the epistemic integrity of the scientific corpus. (Conversely, an upward epistemic adjustment is preferable, even if it can have an impact on the usefulness and productivity of science<sup>27</sup>, because it preserves the reliability of the scientific corpus: it is preferable to wrongly keep scientific issues open rather than to wrongly settle them<sup>28</sup>.)

---

<sup>25</sup>Hansson talks of 'decisional adjustment', but I believe this terminology is slightly misleading, since epistemic decisions are also decisions, as I explained in section 2.2.1.

<sup>26</sup>For example, 'we can treat the evidence for the efficiency of an anti-viral drug as sufficient in one clinical context and insufficient in another, depending for instance on how serious the infection is.' These decisions can be taken without 'changing our answer to the question whether there is sufficient scientific evidence that the drug is effective' (Hansson, 2014, p. 133).

<sup>27</sup>Because it reduces the size of the corpus.

<sup>28</sup>In other words, in science false negatives, which amount to missing an existing phenomenon, are preferable to false positives, which amount to concluding that there is a phenomenon when in fact there is none. Such preference seems indeed to be characteristic of science (scientists are known to be conservative in the inferences they make), and the basis of its authority in society (see e.g. Parascandola, 2010). See section 3.2 for more details.

- upward non-epistemic adjustments, because they contradict the concept of knowledge in general (as sufficiently reliable for practical action), and of scientific knowledge in particular, represented by the corpus (as our most reliable knowledge, apt to any application). (Conversely, a downward non-epistemic adjustment is not problematic in this respect.)<sup>29</sup>

Therefore, these two types of adjustments (downward epistemic adjustment and upward non-epistemic adjustment) are excluded from Hansson’s model. Instead, when there is a discrepancy between epistemic and non-epistemic requirements, one can make:

- either upward epistemic adjustments (if the level required for non-epistemic decision-making is higher): this enables to avoid making upward non-epistemic adjustment<sup>30</sup>;
- or downward non-epistemic adjustment (if the level required for non-epistemic decision-making is lower): this enables to avoid making downward epistemic adjustment. In this case, a separate lower level of evidence for the decision in question is introduced and distinguished from the requirement for corpus entry, so that the integrity of science is preserved.

Now, the epistemic requirement of a claim cannot be adjusted to suit *all* theoretical (2017, p. 219) or practical (2007, p. 266) applications of this claim. The best we can do, according to Hansson (2014, p. 135), is to choose *one* of the levels of evidence suited for one of these applications. In order to preserve the epistemic integrity of science and ensure multipurpose applicability, Hansson proposes to choose the *highest* level of evidence corresponding to the most demanding (theoretical or practical) application. In the end, the final epistemic requirement is the highest of the levels of evidence required by either:

- epistemic entry in the corpus (in which case there is no non-epistemic value influence);
- or non-epistemic theoretical or practical decision-making (in which case there is<sup>31</sup> non-epistemic value influence in the form of an upward epistemic adjustment).

To sum up, either:

- the epistemic requirement is equal to the non-epistemic requirement (figure 2)<sup>32</sup>: no adjustment is needed; there is no non-epistemic value influence<sup>33</sup>;

---

<sup>29</sup>An upward non-epistemic adjustment may lead to raise the epistemic requirement (so that scientific knowledge remains reliable for the envisaged action), thus strengthening the epistemic integrity of the corpus – which is not the case of a downward non-epistemic adjustment. Beside these epistemic considerations, both upward and downward non-epistemic adjustments can be problematic from a practical point of view: for example if someone does not use a vaccine because she requires a higher level of evidence than the epistemic requirement (upward non-epistemic adjustment) or if she uses a drug while requiring a lower level of evidence than the epistemic requirement (downward non-epistemic adjustment).

<sup>30</sup>This is illustrated by Feleppa’s (1981, p. 418) account of the Manhattan project: while initially supposing that theory acceptance took place chronologically before its acceptance ‘in the fuller sense’ (i.e. when scientists announced it ‘to their superiors’ and took into account ‘the social cost factor’, which ‘dr[o]ve the acceptance level upward’), Feleppa then denies that ‘it can be claimed that this contemplation of social cost can be said to be figuring only in some context of application of theory and not, inextricably, in the process of adding hypotheses to theory.’

<sup>31</sup>This is not to say that a non-epistemic theoretical decision always presupposes non-epistemic values (some may not, see footnote (22)), but in such cases where the non-epistemic requirement (due to this non-epistemic theoretical decision) is higher than the epistemic requirement, it does.

<sup>32</sup>This (quite artificial) possibility is not mentioned by Hansson.

<sup>33</sup>Recall that this expression applies to the epistemic requirement (with respect to the incorporation of a claim into the corpus): there is of course always non-epistemic value influence on the *non-epistemic* requirement (with respect to a non-epistemic decision).

- or the epistemic requirement is higher than the non-epistemic requirement (figure 3): a downward non-epistemic adjustment enables us to introduce a separate lower level of evidence for the non-epistemic decision in question, while leaving the epistemic requirement unaffected; there is no non-epistemic value influence;
- or the epistemic requirement is lower than the non-epistemic requirement (figure 4): an upward epistemic adjustment enables us to raise the epistemic requirement to the non-epistemic requirement; there is non-epistemic value influence.




Figure 2: The epistemic requirement is equal to the non-epistemic requirement.

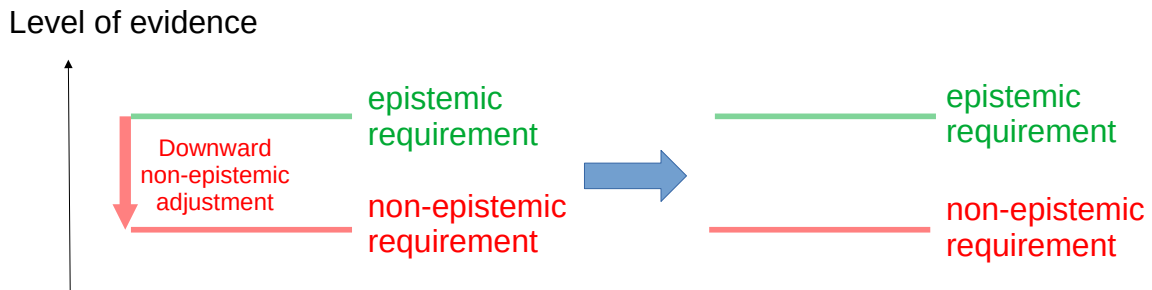


Figure 3: The epistemic requirement is higher than the non-epistemic requirement.

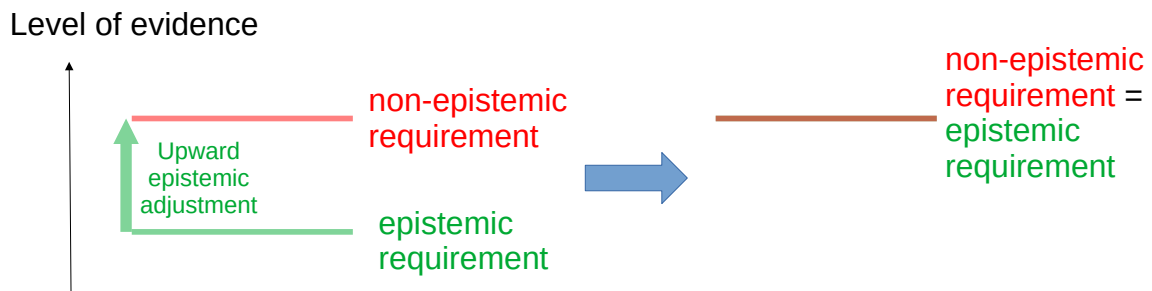


Figure 4: The epistemic requirement is lower than the non-epistemic requirement.

## 2.3 Limits of the model

Note that in the corpus model, values only influence the process of *incorporation* of claims into the corpus (i.e. their acceptance or rejection), not the claims which *already belong* to the corpus. Once incorporated, a claim is considered apt to any application and values are not reconsidered each time a new application is envisaged. Rather, the claim becomes a value-free, purely factual statement, even if the process which led to its acceptance is value-laden (Hansson, 2014, p. 137). This illustrates the process of ‘fixation’ of a factual belief (i.e. what we believe to be a fact) and the fundamental fact-value dichotomy (the need to distinguish our factual beliefs from our other kinds of beliefs: values, preferences, interests, etc.) (Hansson, 2018a, pp. 67, 70). The corpus is a universal repository of factual statements. However, this value-freedom of the corpus only holds as long as new, more demanding applications have not been envisaged: otherwise these applications and the associated values can potentially influence the corpus. If a new application for a given claim is discovered which requires a higher non-epistemic requirement, the epistemic requirement for this claim is raised accordingly. The claim in question, previously incorporated in the corpus on the basis of a lower requirement, is subject to this new requirement, and is excluded from the corpus if it does not meet it. This is not a problem since scientific knowledge is indeed always provisional and the corpus changes constantly (although marginally).

And, of course, in the face of new disconfirming evidence, some claims part of the corpus may subsequently be rejected from it. This aspect, however, is not handled by the corpus model, which only deals with the incorporation of claims into the corpus, not their subsequent *exclusion*. In particular, it does not specify if and how, in the face of new disconfirming evidence, the epistemic requirement is affected, an issue which would deserve further research. Neither does the corpus model explicitly allow for the possibility to (permanently, or for a long time) *suspend judgment* with respect to (the acceptance or rejection of) a claim, even if it allows to do so temporarily, in the sense that it allows to take a non-epistemic decision on the basis of a claim not yet accepted in the corpus (through a downward non-epistemic adjustment). In this respect it is, I think, quite realistic in following scientific practice, where the status of claims cannot be left open indefinitely but must be settled one way or another. This is coherent with Hansson’s rejection of a Bayesian corpus.

## 2.4 Handling of controversial non-epistemic values

According to Hansson (2018a), whereas the philosophy of science has focused on the influence of controversial non-epistemic values in science<sup>34</sup>, most of the non-epistemic values influencing the scientific corpus and having the strongest influence on it, are in fact *not* controversial, such as the assumptions that health is desirable (in medical science), learning is beneficial (in pedagogical science), or a decrease of energy use is advantageous (in technological science). Many of such uncontroversial values are risk-related and require a higher level of evidence than purely epistemic requirements, such as a positive valuation of decreased health risks (in medical science) or increased personal safety (in technological science)<sup>35</sup>. This may explain why<sup>36</sup>, according to Hansson

---

<sup>34</sup>Presumably, such cases are the most talked about in academic philosophy because they call for a solution (even if they represent a minority), contrary to consensual values.

<sup>35</sup>For example, before using a drug, we will require a high level of evidence to show: 1) the efficacy of the drug; 2) the absence of toxicity of the drug.

<sup>36</sup>Although Hansson himself does not relate both claims (neither of which is empirically substantiated in his article), it is plausible to assume that uncontroversial values will more easily influence the production of the scientific corpus than controversial ones.

(2018a, p. 79), the process described by his model of taking the most demanding of epistemic or non-epistemic requirements is implemented continuously and naturally in the scientific corpus (instead of being seen as an external constraint), and corresponds to the reality of scientific practice. It is interesting to note that assumptions about consensual values are also made by proponents of ‘coordinative strategies’, such as Kevin C. Elliott (2013), who presupposes a consensus between scientists, regulators and other stakeholders on the ‘prioritized goals’ (and the corresponding values), e.g. health or economic development, to be pursued in a given research and policy context, and values furthering that goal are then allowed in<sup>37</sup>; or Intemann (2015), whose “aim approach” allows values which promote democratically endorsed aims of research.

Nevertheless, even if it is very difficult to assess as a whole, the influence of non-epistemic values may not be so consensual and straightforward, even for a given context. Indeed, even for values apparently as uncontroversial as promoting health or ensuring safety, there may be different interpretations of the same value (e.g. of what it means to promote health in public health) or conflicts with other values (e.g. safety vs cost-cutting in engineering). Think for example of the different European medicine agencies’ decisions to use the AstraZeneca or the Moderna vaccines against Covid-19 in Europe during the pandemic in 2020-2022. In such cases even values such as promoting health may not be specified in an uncontroversial way: is it the health of the entire population? of a particular subgroup? How do we weigh the risks for specific individuals in comparison to the advantages for the rest of the population? Even in aeronautical engineering – a field where the value of safety is traditionally paramount – values conflicting with safety (such as cost-cutting) may play an important role, by reducing reliability and/or redundancy. For example, at Dassault Aviation company, even for business jets (a field where costs apparently are not the primary concern of customers) the development of the Falcon 5X business jet was made with a strong focus on cost-cutting, contrary to the previous Falcon 7X. This required many trade-offs of safety vs cost requirements<sup>38</sup>. For such conflictual values, approaches which do not presuppose a consensus on values may be more appropriate than Elliott’s or Intemann’s. For example, Carrier (2022) considers values as separate premises or as political commissions and links them to different policy packages implementing them and explicitly stating which values are presupposed (instead of implicitly committing oneself to them)<sup>39</sup>.

Such remarks do not seem, however, to threaten the corpus model. Firstly, the examples given (ambiguity of the value of promoting health, conflict between values of safety and cost efficiency) can be taken to concern the *application* of the scientific corpus (in medicine agencies, or in practical engineering), not the production of the corpus itself. I will return to this question of the frontier between the corpus and its context of application. Secondly, even if there is no consensus on values within corpus production, in principle the procedure of the corpus model can still be applied, if one accepts to ‘pay the price’ of always taking the strictest level of evidence required by all the different values or goals at stake. For example in aeronautical engineering, this would amount to always privileging safety over cost-cutting, even if that would mean more expensive, and potentially less optimised aircrafts). For the AstraZeneca vaccine, the most demanding requirement about safety would lead to ban the use of the vaccine, such as what the Norwegian or Danish medicine agencies

---

<sup>37</sup>This is Elliott’s ‘Multiple Goals Criterion’ (2013, p. 381) (see footnote (43)). Elliott is aware that such a consensus is a ‘challenge’, which he nevertheless assumes as resolved (on a case-by-case basis, for each ‘particular context’) for his model to apply.

<sup>38</sup>This case is taken from my experience as a thermal systems engineer in this company for four years.

<sup>39</sup>According to Carrier, this conditionalization of scientific advice according to explicitly presupposed values enables to keep, to a large extent, the value-free ideal, by providing ‘*alternative value-laden policy packages* which combine facts, scientific accounts and non-epistemic premises’ (2022, p. 16).

did<sup>40</sup>. Nevertheless, in practice taking the maximum requirement can be challenging, as we shall see in section §3.

## 2.5 Advantages of the corpus model

There are several advantages to the corpus model. Firstly, the model obviously has the merit of focusing on the scientific corpus, i.e. on the scientific claims (and the associated required level of evidence), and not on more evanescent matters like values themselves (which are treated indistinctly by the model), social settings of the scientific community, expectations of other stakeholders, etc. Of course, values are, and must be, taken into account by the model, since the model cannot avoid the need to make value judgments at some point, namely for the choice of the level of evidence required for a claim in view of its applications. But this seems inevitable for any model of value influence.

By providing a clear and rigorous procedure to be followed, the model is operational and easy to implement<sup>41</sup>. Once the epistemic and non-epistemic requirements are set, the definite procedure prevents any ambiguity or latitude in personal choices, and associated danger of bias. Moreover, this procedure is fairly *simple*, and, most importantly, it is always the same, regardless of the claim, values or goals considered – in other words, the model is *universal*, and this universality further contributes to the simplicity of the model. The model is also *systematic*, in the sense that it distinguishes the various possible cases and what to do in each. Finally – and this may be its biggest advantage – the corpus model systematically preserves (or even strengthens) the *epistemic integrity* of science, since the final corpus entry requirements correspond to the maximum of the epistemic and non-epistemic requirements. Incidentally, the model also preserves the *unity* of science, since it relies on a unique, interdisciplinary corpus.

As far as I can see, all these features (clarity and rigour, simplicity, universality, systematicity, preservation of epistemic integrity) favour the corpus model over other models in the literature which are less clear and systematic, leave margins of interpretation, work on a case-by-case basis – like those depending on specific values (e.g. Kourany, 2010)<sup>42</sup> or goals (e.g. Kevin C. Elliott, 2013)<sup>43</sup>, combine different of these approaches (Kevin C. Elliott, 2022) and/or do not systematically guarantee the epistemic integrity of science. What is more, such features are surely beneficial for public trust in the way values may influence science<sup>44</sup>. As far as I know, Hansson’s model is the

---

<sup>40</sup>Given that other vaccines were available in these countries, such measures do not seem to have had any health effects. Nevertheless, since there was a global lack of vaccines, and that AstraZeneca vaccines were a much better option than no vaccine at all, it would have been problematic to just throw them away, or to send them to poor countries (where they could contribute to increase vaccine hesitancy, which is already higher than in rich countries, by providing populations there with a vaccine which rich countries do not want for themselves).

<sup>41</sup>At least for the *users* of the scientific corpus, i.e. other scientists who use it inside science to produce further scientific knowledge, people (including decision-makers) who apply it outside science to a variety of purposes, and the general public at large who gets a simpler view of science. On the other hand, for the *producers* of the scientific corpus, i.e. the scientists who have to apply the corpus model, this can be a quite complicated task, since it requires to identify all reasonably foreseeable applications (more on this in section 3.1). Of course the VFI is even simpler since it does not at all take into account non-epistemic values, but it is excluded from our endeavour here, which purports to find a VLI model.

<sup>42</sup>For whom the values allowed are those that ensure ‘human flourishing, what makes for a good society’ (2010, p. 68).

<sup>43</sup>Elliott’s ‘multiple goal criterion’ makes values depend on the goals of scientific knowledge production: according to this criterion, ‘[a] particular value can appropriately influence a scientist’s reasoning in a particular context only to the extent that the value advances the goals that are prioritized in that context’ (2013, p. 381).

<sup>44</sup>Of course, value influence on an case-by-case basis should not be ruled out *a priori*. In principle at least there



only one to have such features. Surprisingly, it is largely ignored by the rest of the literature on values in science. This is unfortunate, since a confrontation between this model and other ones (which mainly work on a case-by-case basis) may be fruitful, and potentially enable progress on the issue of value influence. In any case, it is worthwhile to note that the corpus model already answers some concerns voiced in the literature. Of course, the following are only examples, and it would take another article (at least) to confront more systematically the corpus model to others in the literature.

For example the corpus model answers concerns about the various senses of what it means to accept a scientific hypothesis or theory, a critique formulated against Douglas’s (2009) account of value influence (Morgan 2010, p. 424; John, 2012, p. 219). For instance, Morgan (2010, p. 424) criticises Douglas for failing to identify the ambiguity in the expression ‘theory acceptance’, which can mean: “accept as a true belief”, “accept as a theory on which to base policy”, “accept as a theory on which to develop technology”, “accept as a theory worthy of further pursuit”, etc. Different meanings incorporate values differently and pragmatic arguments that work for one meaning might not work for another.’ We see that the corpus model, through the above delineated procedure, enables to distinguish between various senses of theory acceptance: a theory accepted as a true belief is incorporated in the corpus; for the other cases (accept a theory on which to base policy, to develop technology or to pursue further) the theory can either be incorporated in the corpus or not, according to the situation (discrepancy or not between epistemic and non-epistemic requirements, and in the first case upward epistemic adjustment or downward non-epistemic adjustment). As we have seen, it may happen for example that the theory is not accepted in the corpus while still serving as a basis for policy making. Or, conversely, we may need, for policy making purposes, to submit the theory to criteria of evidence even stricter than normal epistemic requirements. Similarly, the corpus model takes into account the possibility to either accept a claim, reject it, or withhold judgment while still recommending to take action as if the claim were true (John, 2012, p. 219). In the same way, the model handles Lacey’s (2017) distinction between impartially holding a claim (‘meaning, roughly, that the claim is so firmly established as to not require further testing’), adopting it (i.e. ‘treat[ing it] as a basis of further research’) or endorsing it (i.e. ‘tak[ing] it as a basis for decision-making’) (K. Elliott and Steel, 2017, p. 4).

More generally, Holman and Wilholt (2022) divide the various demarcation strategies for identifying legitimate value influence into five types: axiological (based on the values themselves), functionalist (based on the role values play), consequentialist (based on the concrete consequences that the adoption of values has on achieving certain aims), coordinative (according to whether certain value choices correspond to the expectations of some stakeholders), and systemic (dependent on the ‘social set-up’ within which research takes place) strategies. The corpus model can be classified as a functionalist strategy: non-epistemic value influence is accepted only if it contributes to increase the level of evidence required for entry into the corpus. Moreover, the corpus model also answers the three concerns historically associated with the VFI identified by Holman and Wilholt (2022, p. 214), namely veracity, universality and authority<sup>45</sup>: it preserves or even strengthens the

---

is nothing wrong with this option (one does not necessarily need a general law, and in the absence of it one can for example try to identify case laws on the basis of case studies). The impact of value influence on a case-by-case basis on the general public opinion is also not clear (I am not aware of any study investigating this particular point, in any case). Nevertheless, there is something appealing in having a universal, simple model, and one can assume it would have a positive impact on the general public’s trust. Conversely, one can assume that a case-by-case model might spur suspicion, because having a set of rules for each case is more complicated as a whole (even if each rule is simpler than the corpus model) and because the epistemic requirement is adjusted to each case (hence possibly entertaining the suspicion of partisan science).

<sup>45</sup>Although this is not the place to argue for it, I believe universality and authoritativeness stem from veracity.

epistemic integrity of scientific knowledge (veracity); it ensures that this knowledge is apt to any use, whether theoretical or practical (universality); and presumably the knowledge thus produced is more trustworthy (authority). Therefore the corpus model answers Holman and Wilholt's (2022, p. 219) own demarcation criteria, and seems to be a very good candidate model for value influence.

However, all the advantages previously mentioned come at a price. Indeed, the universality, systematicity and preservation of epistemic integrity (through taking the maximum requirement) result in a certain *rigidity* of the model, which is not optimised for applications less demanding than the most stringent one for a given claim. It also presupposes that this most stringent application can be identified in the first place. These difficulties are discussed in section §3.

### 3 Critical discussion

In this section I review some difficulties bearing on the characteristic feature of the corpus model which consists in taking the maximum of the epistemic or non-epistemic requirements, i.e. the maximum level of evidence required for entry into the corpus or for theoretical or practical decision-making:

1. the difficulty to identify this most demanding requirement (maximum level of evidence) (§3.1);
2. the fact that this maximum requirement is not optimal with respect to all other (epistemic or non-epistemic) requirements (§3.2);
3. the fact that this maximum taking proceeds from a potentially detrimental intra-scientific preference for false negatives over false positives (§3.3);
4. the issue of whether there can nevertheless be requirements higher than this maximum (in other words whether a claim belonging to the corpus could still be deemed not reliable enough for some applications) (§3.4).

While the first issue has to do with the applicability of the corpus model<sup>46</sup>, the second and last one have to do with its improvement and re-conceptualisation, and the third one with a presupposition on which it relies.

#### 3.1 Difficulty to identify the most demanding requirement

First, the question naturally arises as to whether it is *possible*, for a given claim, to identify this maximum requirement, in other words to identify *all* the (theoretical or practical) applications of this claim, in order then to take the most demanding one. The impossibility to foresee all future applications of one's research is one of the arguments used by Polanyi (1962/2000) in favour of scientists' exemption of moral responsibility, in other words against any influence of (moral) values. Polanyi takes the example of the atomic bomb, an application of the theory of relativity which took place 40 years after its publication, and which Einstein could not possibly have anticipated. Since

---

Following their study of various demarcation strategies, Holman and Wilholt (2022, p. 219) acknowledge that universality and authoritativeness are not sufficient without veracity.

<sup>46</sup>Note that I take for granted here the normative stance – consistent with the VLI – according to which scientists *should* be held responsible for the potential applications of their claims.

scientists are often unable to foresee both the results and the practical applications of their research, they cannot be held accountable of the potential detrimental applications of their research, according to Polanyi. To some extent, the argument of impossibility (or rather, ‘impracticability’, 149) of anticipating the future applications of one’s research had already been made by Bridgman (1947). The latter does not properly claim that it is impossible, but rather lists several arguments against it (although to some extent, he also seems to accept such a possibility), such as the special responsibility it would confer to scientists, in contradistinction to other members of society (but for a rebuttal of this claim, see Douglas, 2009, 68–75, who argues that this special responsibility proceed from the general responsibility all members of society have); the fact that it would go against the division of labour foundational of our modern societies; that it would be counter-productive because it would force scientists to do things which they are not good at (foreseeing the potential uses of their discoveries), while at the same time hindering them from doing things which they are good at (i.e. research); that it is the industry, where the applications are made, which should control these applications; that there are other ways for society to deal, as a whole, with the potential detrimental applications of claims, which in fact will disappear by themselves in a good society; and that if really society wants scientists to be responsible, it should create ‘mechanisms’ and ‘opportunities’ for those scientists specially attracted to this task to be able to check the applications. Bridgman also rightfully notes that scientists’ values need not reflect those of society (for an answer to this point, see e.g. Intemann, 2015, where values are legitimate when they promote democratically endorsed epistemological and social aims of research).

Of course, the scientist cannot be asked to take into account, and (retrospectively) be held responsible for, *any future, unforeseeable* applications of her research. In the corpus model, however, the idea seems to be that we are only concerned with the currently *known* (and, we might add following Douglas (2009, pp. 66–86), *reasonably foreseeable*) applications of a claim. Thus, in principle, the corpus model should be applicable. And if new applications are envisaged or discovered, corresponding new standards of evidence are required and can influence the corpus (see section 2.3). However, in practice, identifying all the (known, not to speak of reasonably foreseeable) theoretical or practical applications of a claim may be difficult for a single scientist (or even a given scientific community), especially for a theory or hypothesis of high generality (such as law-like hypotheses) and/or which has many applications in various fields<sup>47</sup> (such as a theory or hypothesis in a fundamental discipline like fluid mechanics, which has theoretical applications in many others disciplines<sup>48</sup> or practical applications in many different industries<sup>49</sup>). Of course, one can object with Douglas (2009, p. 83) that one should only anticipate the ‘reasonably foreseeable’ applications<sup>50</sup>. But the counter objection would be that this notion of ‘reasonably foreseeable’ is difficult to define precisely, and may be unrealistic<sup>51</sup>. For theoretical applications, it seems unlikely

---

<sup>47</sup>A remark already made by Jeffrey (1956, p. 242) (even if I do not subscribe to his conclusion that the scientist should therefore refrain from accepting or rejecting hypotheses, and instead only provide probabilities for them).

<sup>48</sup>For example mechanical, civil, chemical and biomedical engineering, geophysics, oceanography, meteorology, astrophysics, or biology.

<sup>49</sup>Such as the car industry, aeronautical industry, naval industry, etc.

<sup>50</sup>Indeed, Douglas does not hold the (implausible) claim that scientists are responsible for any application of their research, only for those that they could reasonably have foreseen. Note that Douglas’s claim is not explicitly related to Hansson’s concept of multi-purpose corpus.

<sup>51</sup>One reviewer suggested that work in ethics of expertise (such as Elliott’s (2006)) might be helpful for framing a notion of scientists’ ‘reasonable’ effort. Although I agree Elliott’s transposition of the notion of informed consent from biomedical ethics to ethics of expertise brings insight, for example for conceptualising the notion of ‘reasonableness’ (with respect to the ‘reasonable person standard’ of informed consent (Beauchamp and Childress, 1994/2012, p. 126), i.e. the information that a reasonable person would want to receive when faced with a *particular*

that a scientist (by definition specialised in her discipline) may master all the potential applications of her claim in other disciplines (in spite of the growing interdisciplinarity mentioned by Hansson). For practical applications, the scientist *qua* expert, specialised in the practical applications of her research, is surely in a better position for doing so than the scientist *qua* researcher, who contributes to the production of the scientific corpus. But apart from the fact that these two roles may be more or less separated<sup>52</sup>, not all scientists have expert activities. Even for an expert, the variety of potential practical applications, exceeding the expert’s area of specialisation, may represent a serious difficulty.

In sum, the applicability of the corpus model is somewhat put into question with respect to the identification of the maximum requirement: even if the latter is *in principle* simple to implement, *in practice* it can be challenging. The cognitive economy allowed by the multi-purpose scientific corpus is obtained only at the price of a cognitive burden paid by the scientists incorporating the claims into the corpus, who have to be somewhat omniscients. In other words, while the corpus *concept* is cognitively economical, the corpus *model* is cognitively demanding. It is therefore legitimate to ask if multiple, application-adjusted corpora are not more realistic cognitively speaking, since they require a less far-reaching knowledge of all possible applications. I come back to this suggestion at the end of the next subsection.

### 3.2 Non-optimality with respect to all other requirements.

Even if taking the maximum level of evidence required by all the applications of a claim is possible, this means that this requirement is too high with respect to all other applications, in other words that the corpus is not application-optimised. This is clearly a drawback due to the rigidity of the corpus model, which has to do with its universality and uniqueness. Taking the most demanding requirement can in principle have consequences both on the theoretical and practical levels:

- On the theoretical level, by reducing the size of the scientific corpus, i.e. the number of accepted claims: this drawback does not affect the epistemic integrity of the corpus, but it can have an impact on the productivity of science (by reducing the number of available results on which future scientific knowledge will be based) (Hansson, 2018a, pp. 68–69, 79–80). Nevertheless, according to Hansson this is a price we are willing to pay in order to ensure the epistemic integrity of science, all the more so since lowering entry requirements would also have an impact on the productivity of science (by jeopardising future research on the basis of unreliable results), in addition to threatening the epistemic integrity of the corpus. According to Hansson, it is preferable to have too high rather than too low entry requirements into the corpus: this presupposition will be discussed in the next subsection.
- On the practical level, we can add that reducing the size of the corpus may similarly reduce its (practical) usability for other, less demanding applications. But an objection similar to Hansson’s (in the previous bullet point) would be that lowering entry requirements may also reduce the practical usability of the corpus by reducing its reliability.
- What is more, for both theoretical and practical applications, Hansson’s model enables (theoretical and practical) downward non-epistemic adjustments to apparently ‘bypass’ these

---

decision), he always ‘assum[es] that one can at least roughly identify potential decisions that might be affected by the information’ in question (643). But this is precisely my point: that these decisions cannot always be identified – even if, ‘in many cases, it may be fairly obvious to scientific experts that their disseminated information is likely to be used for particular decisions’ (649).

<sup>52</sup>I do not enter this debate here and refer to Gundersen (2018).

drawbacks (by deciding for example to pursue research or taking a medical treatment on the basis of a yet-unaccepted hypothesis), while keeping the epistemic integrity of the corpus untouched.

Let us now see how this works in practice. The objection against taking the most demanding requirement for entry into the corpus was in fact already made by Jeffrey (1956, p. 245), who took the example of the hypothesis that a given polio vaccine is free from virulent virus, with obviously different safety requirements for inoculation to humans or to monkeys: ‘[...] the scientist should refrain from accepting or rejecting hypotheses since he cannot possibly do so in such a way as to optimise every decision which may be made on the basis of those hypotheses. *We note that this difficulty cannot be avoided by making acceptance relative to the most stringent possible set of utilities (even if there were some way of determining what that is) because then the choice would be wrong for all less stringent sets.* One cannot, by accepting or rejecting the hypothesis about the polio vaccine, do justice both to the problem of the physician who is trying to decide whether to inoculate a child, and the veterinarian who has a similar problem about a monkey.’ (245, my italics) Jeffrey recommends instead that the scientist: either provides the probability of the hypothesis under study (in which case everyone can make their own decision based on it and their particular problem at hand)<sup>53</sup>; or ‘takes on the job’ of making a separate decision of accepting or rejecting the hypothesis in each case (instead of accepting or rejecting it ‘once and for all’). The latter possibility represents, in other words, multiple, application-adjusted corpora. To use Jeffrey’s example, instead of accepting or rejecting absolute claims like ‘this vaccine is free from virulent polio vaccine’, the scientist would accept or reject claims like ‘this vaccine is free from virulent polio virus for inoculation of a human,’ or ‘this vaccine is free from virulent polio virus for inoculation of a monkey’ (for a given vaccine, the first claim could for example be rejected and the second accepted).

This latter view agrees in fact to some extent with Hansson’s (2007) prior VFI conception, according to which pure science should be value-free while applied science can be value-laden (see section 2.1.2). There, while Hansson rejects *multiple* corpora in *pure* science because of the growing interdisciplinarity between disciplines, he also rejects a *unique* corpus in *applied* science because of the variety of practical uses of science<sup>54</sup>. Indeed, Hansson explicitly rules out the possibility of a (unique) ‘[practical] application-adjusted corpus’, whose criteria of evidence would ‘suit the practical purposes for which we intend to use it’, because ‘practical purposes differ so that there are no stable and general “practical” standards of evidence’ (2007, p. 266):

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

---

<sup>53</sup>One of the reviewers remarked that this is what the Intergovernmental Panel on Climate Change does in its reports for decision-makers. This solution seems indeed suited to the nature and purpose of these reports, which differ from the scientific corpus in several respects (e.g. they are co-constructed with governments), and are intended to provide decision-makers with an overview of the state of current scientific knowledge about climate change in order to help them devise policy. However, it seems ill-suited to the practice of science itself, for the reasons advanced by Hansson (cognitive unmanageability of a corpus made of probabilistic statements).

<sup>54</sup>In both cases, Hansson would presumably prefer a unique scientific corpus, corresponding to his conception (see section 2.1.2). But whereas the situation in pure science allegedly corresponds to his conception, it is not the case in applied science, at least as long as he has not come up with his corpus model (compare to footnote (17)). In the latter Hansson treats theoretical and practical applications on a par.

by an airplane constructor or by someone designing a decoration. (Hansson, 2007a, p. 266)

Thus, because of the variety of our practical uses of science, there is no unique way to adjust the standards of evidence to all these practical uses. For example, we will require a higher level of evidence for establishing the absence of toxicity of a drug against a minor disease than we would against a life-threatening condition (for which we will accept a greater uncertainty about side effects).

But from 2014, the corpus model enables to manage this almost infinite variability of our theoretical and practical uses of the corpus, through 1) taking for each claim the maximum requirement of all its applications, in order to ensure the integrity of the corpus as well as its universal applicability; 2) offering the flexibility, with downward non-epistemic adjustments, to nevertheless ‘accept’ claims for non-epistemic (theoretical or practical) purposes. Thus, even if the absence of toxicity is not established in the most demanding case (i.e. against a minor disease), we may nevertheless use the drug – following a downward non-epistemic adjustment – against a life-threatening condition, because we deem the benefit-risk ratio favourable (i.e. the therapeutic effects are more probable than the secondary effects). Similarly, we will require a higher level of evidence for establishing the safety of a vaccine in normal conditions than in conditions of emergency, for which, even if the safety of the vaccine is not established, we may nevertheless use it following a downward non-epistemic adjustment. We will require a higher level of evidence for a claim about the strength of a material for building an airplane than for building a decoration (for which, obviously, we can use the material following a downward non-epistemic adjustment).

In general, we can identify the following cases for claims to be handled by the corpus model:

- for a ‘positive’ claim (stating a desired property or state, such as, in pharmacology, the absence of toxicity, the safety or the effectiveness of a drug):
  - practical applications may require higher levels of evidence than incorporation into the corpus (in which case, according to the corpus model, the epistemic requirement will be raised accordingly)
  - a preference for false negatives over false positives means to prefer running the risk of missing out a drug with beneficial effect over running the risk of using a drug with no effect: such a preference may be justified both for theoretical and practical applications;
  - downward non-epistemic adjustments (e.g. in case of serious sickness or epidemic for example) enable to alleviate or suppress the risk associated with this preference .
- for a ‘negative’ claim (stating an undesired property, such as the presence of undesirable side effects of a drug):
  - incorporation into the corpus requires higher level of evidence than practical applications (for example, we will not wait that a danger is scientifically confirmed before trying to avoid it);
  - a preference for false negatives over false positives means to prefer running the risk of accepting a drug with detrimental side effects over running the risk of not using an effective drug because of exaggerated negative effects: such a preference is generally not justified for practical applications (although it may be justified for theoretical applications);

- downward non-epistemic adjustments enable to alleviate or suppress the risk associated with this preference.

Thus, if we subscribe to Hansson’s arguments for cognitive economy and corpus manageability against both a Bayesian corpus and multiple application-adjusted corpora (see section 2.1.3), both of Jeffrey’s options seem unrealistic. On the contrary, for Hansson we need a provisionally fixed (i.e. a claim is either rejected or accepted, not provided with probabilities) and a universal, multi-purpose (i.e. a claim is not accepted or rejected according to the application) corpus.

While I agree with Hansson’s rejection of a Bayesian corpus as unrealistic, I think that his rejection of multiple application-adjusted corpora may perhaps be challenged. Of course, such application adjustments would not concern hypotheses of high generality, such as claims regarding the chemical structure of DNA (as in Hansson’s (2018, p. 71) example), or law-like hypotheses (like, say, Maxwell’s equations). Rather, they seem plausible for hypotheses of low generality in certain disciplinary fields, or for recently accepted hypotheses which are subject to more uncertainty than older hypotheses of the corpus. For example, regarding claims like those concerning the safety or efficacy of a drug or vaccine, we may conceive different corpus entry requirements adjusted to the applications – in other words, different application-adjusted corpora. Thus, the safety or efficacy of a drug would not be established ‘absolutely’, but only relatively to the clinical context – as, indeed, seems to be the case in actual medical practice, if I am not mistaken. An absolute statement like ‘this drug is safe (or efficient) for treating this condition’, or ‘this vaccine is free from virulent polio virus’ (to take Jeffrey’s example) would not be sufficient, but would rather be related to the context of use (e.g. ‘this vaccine is free from virulent polio virus for inoculation to humans’), in order to evaluate the risk benefit ratio of the therapeutic vs adverse effects of the drug. A drug can be deemed safe, or efficient, in one context, because its beneficial effects are greater than the adverse ones, but not in another where this is not the case. This is indeed the kind of information that the indications and usage section of the prescribing information of a drug provides, in order to serve as guidelines for practitioners and patients. Of course, it can be argued, as Hansson (2014) does, that using a drug in a clinical context is ultimately a matter of practical decision. Nevertheless, even if the indications and usage of a drug do not strictly belong to the corpus (but rather to its practical application, as guidelines for practitioners and patients), they are supposed to stem from the best available literature on the subject.

In the end, the issue amounts to where one places the frontier between the claim to be incorporated in the corpus, and its context of application. The decision of where to set the frontier determines if there is one corpus or multiple corpora, and also influences which model of values in science is relevant (for example, the corpus model does not work with multiple corpora). Conversely, the choice of a model for values in science influences where the frontier is set (for example, the corpus model excludes the mention of each application in applied science).

### 3.3 Questioning the intra-scientific preference for false negatives over false positives

As we saw in the previous subsection, Hansson prefers too high corpus entry requirements (which threaten only science’s productivity) rather than too low (which threaten both science’s integrity and productivity). In other words it is preferable to miss an existing phenomenon (false negative) rather than to conclude that there is a phenomenon when there is in fact none (false positive)<sup>55</sup>.

---

<sup>55</sup>See footnote (28).

The corpus model relies on the assumption that within science, false negatives are always preferable to false positives, in conformity with the traditional conservatism and cautiousness of scientists<sup>56</sup> (this view is expressed e.g. in Hansson, 2008b, p. 7). This presupposition, which is part of Hansson’s definition of science and of the corpus concept, ensures that the integrity of science is preserved, and seems conceptually prior to the corpus model. It is this presupposition which apparently leads Hansson to take the most demanding requirement (the highest level of evidence required by epistemic considerations, or theoretical or practical applications) in order to ensure that the integrity of science is preserved. Hansson recognises however that outside science (i.e. in practical applications) false negatives may have ‘much more serious’ consequences (2008, p. 7), and therefore are not necessarily to be preferred in this context. This danger is supposed to be avoided by downward non-epistemic adjustments, which should be based on any relevant scientific data which has not yet given rise to scientific publication<sup>57</sup>.

However, not everybody agrees that false negatives should systematically be preferred to false positives within science, in other words that error avoidance is preferable to truth attainment. For example according to Parascandola (2010), such a conservative attitude runs counter to science’s aim of increasing knowledge and reducing uncertainty. A strong aversion to risk may lead scientists to the paradoxical position of gathering less evidence (because each new experiment introduces additional risks of inductive error if the investigator uses the new evidence produced to draw new inferences), hence countering science’s aim of understanding the world and getting closer to the truth. According to Parascandola some degree of error is inevitable in science, and as the scientific endeavour progresses, it is normal that the number of false findings increases, but so does the number of true ones. Thus, there is no general rule for deciding which type of error is preferable, sometimes false negatives might be worse than false positives, and trade-offs should be made on a case-by-case basis. However, I think Hansson’s point of view precisely answers Parascandola’s concern about hindering the progress of science: as Hansson claims, accepting false claims (i.e. false positives) would precisely hinder the progress of science by leading it into all sorts of blind alleys.

Consider for example Hansson’s example of a claim that a drug has some undesirable side-effects (2020, p. 383). It seems that here a false positive (believing wrongly that the drug has side-effects, with the risk of not using a drug which is in fact safe) would be preferable to a false negative (believing wrongly that the drug has no side-effects, with the risk of using it and putting the patients at risk). But this case does not invalidate Hansson’s general assumption that false negatives are preferable to false positives *in science*, since what is at stake here is not the acceptance of the hypothesis in the corpus, but a practical decision (made by clinical practitioners, or decision-makers such as medicine agencies). This decision can be made while still avoiding incorporating a false claim into the corpus: in case of doubt about this claim, a downward non-epistemic adjustment can be made (we require less proof for acting as if this claim were true, than we would for accepting it into the corpus), in accordance with the corpus model. This illustrates precautionary approaches like the principle of precaution (according to which ‘where there are

---

<sup>56</sup>The preference for false negatives over false positives is one illustration of core scientific values such as conservatism, cautiousness, restraint, dispassion, skepticism (see e.g. Brysse et al., 2013, p. 334, who identify under-prediction – rather than over-prediction – as another consequence of such values). The purpose of the present article is not to conceptualise an *ethos* of science like Merton (e.g. 1938), but it is safe to say that these values are very consensual in the scientific community.

<sup>57</sup>Hansson insists that such non-epistemic adjustments should nevertheless be based on the same criteria as those for entry into the scientific corpus (namely the same type of evidence should be taken into account, and the same assessment should be made of how strong the evidence is, as for incorporation into the corpus), except regarding the required level of evidence (which can be lower than for regular incorporation into the corpus) (2008, pp. 145–146).



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’ (UN, 1992)) and the principle of *primum non nocere* (‘first, do not harm’, according to which non-intervention is preferable to intervention if the latter can cause harm to the patient (see Hansson, 2020b, p. 386)). Both principles can be interpreted, in the face of a potential danger (expressed in the form of a claim stating a harm), as a preference for false positives (wrongly accepting this claim) rather than false negatives (wrongly rejecting it). But again, both principles have to do with practical action, not acceptance into the corpus<sup>58</sup>.

Nevertheless, there are other cases for which the systematic preference for false negatives over false positives *in science* (characteristic of scientific conservatism), on which the corpus model is based, might be problematic. Although I think Hansson is essentially right to argue for the intra-scientific preference of false negatives over false positives, here I provide two examples which challenge this general rule (on the theoretical and practical levels respectively).

### **3.3.1 Scientific conservatism can reinforce extra-scientific values detrimental from the epistemic point of view**

Lewandowsky et al. (2015) provide a general reflection and a case study of how influence of extra-scientific values (which they deem ‘seepage’), reinforced by intra-scientific values, can threaten the integrity of the scientific corpus. More precisely, they show how the pervasive climate denialism of society to which U.S. climate scientists are continuously exposed (although they do not subscribe to it themselves) can have a detrimental effect on scientific peer-reviewed literature, by: making scientists take seriously an alleged ‘pause’ in global warming, which is in fact neither backed by evidence nor consistent with standard scientific practice; as well as put into doubt, and require a higher level of evidence for, already accepted claims about global warming. Although the case is mainly about detrimental extra-scientific value influence (which in itself does not threaten the corpus model), nevertheless it shows<sup>59</sup> that intra-scientific values of ‘conservatism’, ‘reticence’ and “‘erring on the side of least drama”’ i.e. a tendency to systematic under-predictions (in a reference to Brysse et al., 2013, see below) can strengthen this detrimental influence.

### **3.3.2 Scientific conservatism can lead to bad societal consequences**

Scientific cautiousness and conservatism are the core subject of Brysse et al.’s (2013) study, which shows that such values can actually lead to under-predictions of key attributes of global warming for policy purposes. More precisely, they show that core scientific values such as ‘caution’, ‘skepticism’, ‘dispassion’, ‘restraint’, ‘moderation’, make scientists err on the side of less rather than more alarming predictions, and require higher levels of evidence for surprising, dramatic or alarming conclusions than for conclusions less surprising, less alarming, or more in line with the scientific status quo (a tendency which they dub ‘erring on the side of least drama’ or ‘ESLD’). Such ESLD is documented in assessment reports of the Intergovernmental Panel on Climate Change (IPCC), as well as NASA press releases, where key attributes of global warming (such as mean temperature

---

<sup>58</sup>So that Parascandola’s (2010) (rather confused) use of these principles as illustrative of his claim (that avoiding type I errors in science is not always preferable) is not vindicated, because they do not concern incorporation into the corpus but practical action.

<sup>59</sup>I find this aspect of the paper less convincing and more programmatic (whereas I find the main argument of extra-scientific ‘seepage’ very compelling). On the other hand, Brysse et al. (2013)’s paper is centred on this aspect (intra-scientific values of conservatism and cautiousness) and is very compelling in this respect, but it is about IPCC assessment reports (and not regular peer-reviewed literature).

change, sea level rise, and atmospheric carbon dioxide concentration) are underestimated. It thus has to do with practical applications of the corpus (for policy purposes). Now even if such ESLD is not in principle a threat to the corpus model (which enables to consider directly the available data and ‘bypass’ the scientific corpus through downward non-epistemic adjustments, and therefore enables to take precautionary approaches and avoid *practical* systematic preference for false negatives over false positives), in practice this is not done in the case of climate change, where IPCC assessment reports precisely represent the basis (as the relevant summary of scientific literature and data) on which policy decisions are to be taken. This case thus illustrates how scientific cautiousness (including systematic preference of false negatives over false positives) can have detrimental social consequences.

### 3.3.3 Probing the corpus model

These two cases suggest that scientific conservatism (which includes a systematic preference for false negatives over false positives) can have detrimental consequences both on the theoretical and practical levels (for the content of the corpus as well as for its use for policy decisions). Consequently, the core idea, in the corpus model, of systematically taking the maximum requirement (which comes from a systematic intra-scientific preference for false negatives over false positives), is threatened – and so is the corpus model itself, since it cannot be applied systematically and universally. Such counter-examples (and especially the first, less-convincing one, see footnote (59)) are clearly not sufficient to discard the model altogether. Nevertheless, they cast doubt on its universal applicability, and call for further studies to better assess this phenomenon and its extent, in order to be able to better assess the relevance of the corpus model itself, and possibly modify it. In particular, we might want to reconsider multiple, application-adjusted corpora instead of one multi-application corpus<sup>60</sup>.

## 3.4 Beyond the maximum requirement

### 3.4.1 Can a claim in the corpus still be considered not reliable enough for practical application?

As we have seen, one important idealisation of the corpus model is that it posits a unique, all-purpose scientific corpus, whereas we could imagine several corpora according to the envisaged (theoretical or practical) application. According to Hansson (2018a, pp. 77–78), once a claim has been incorporated in the corpus, it should be apt to any application, otherwise this would contradict the concept of knowledge in general, and of scientific knowledge in particular, as our universal, most reliable knowledge, apt to any application. Hence, Hansson excludes the possibility according to which a claim, once incorporated in the corpus, would still need a higher level of evidence for some specific practical decision (typically, for an application requiring a particularly high level of confidence). In such a case, his model envisions, as we have seen, an upward epistemic adjustment corresponding to this higher non-epistemic requirement, and the final entry requirement into the corpus corresponds to this higher non-epistemic requirement.

What is more, according to Hansson (2018, 79, my translation) such a situation of discrepancy between epistemic and non-epistemic requirements should normally not even happen in scientific

---

<sup>60</sup>If we have to choose a unique, universal rule, then Hansson is probably correct to defend an intra-scientific preference for false negatives over false positives. The question is rather whether we should aim at a universal rule in the first place.

practice, since it is ‘anticipated’ by the process of scientific knowledge production, whereby epistemic requirements are raised to the non-epistemic requirements: ‘The corpus is continuously adjusted not only to the needs of scientific research but also to the practical needs which have stricter evidence requirements than the scientific process itself. If such a divergence [between epistemic and non-epistemic requirements] should happen, it would generally be treated through an adjustment of entry conditions into the corpus rather than through special arrangements for the practical decisions in question [i.e. higher non-epistemic than epistemic requirements].’ Here Hansson’s claim seems to be descriptive of scientific practice. As such, it should be able to be empirically confirmed or disconfirmed.

However, it seems *prima facie* plausible to think of situations for which claims taken to be scientifically established could still be deemed as insufficiently reliable for some specific theoretical or practical application(s). Such a possibility might for example arise in fields requiring especially high safety levels, such as engineering (for instance aeronautical engineering). Or think of a new research program in particle physics or astrophysics requiring to build very expensive equipments (such as particle accelerators or telescopes) on the basis of previous research. For such theoretical or practical applications, we might want to require higher levels of evidence than those commonly accepted in the scientific literature. Although the corpus model is designed to take care of such discrepancies (by raising the epistemic requirement), if in actual practice the corpus is not used as the model envisions, a revision of the model may legitimately be envisaged.

### 3.4.2 Theoretical and practical engineering

Before studying some examples in engineering to illustrate the above-described discrepancy situation, it will be helpful to first define engineering itself. Following Hansson (2007; 2013) (for whom it seems to be implicit), I take (theoretical, or scientific) engineering (or engineering science) to be synonymous with technology (or technological science). It is a ‘science in the large sense’ contributing to the scientific corpus (see my remarks above) (Hansson, 2018a, pp. 60–65). Technological science uses scientific methodology to investigate human-made objects. It must be distinguished from two similar endeavours: (theoretical) natural science, which uses scientific methodology to investigate natural objects; as well as from applied natural science, which uses *results* from (theoretical) natural science (e.g. available formulas and theories) to investigate human-made objects, among other things<sup>61</sup> (Hansson, 2013, pp. 17–18, 22; Hansson, 2007b). In my opinion, engineering can (conceptually, at least) be divided into:

- theoretical or scientific engineering (or engineering science, or technological science, or technology), which is an applied science (along e.g. medical science) producing knowledge incorporated in the corpus;
- and practical engineering, where in general no new knowledge is created, but rather where the findings of theoretical engineering are used (in other words, it is an application of theoretical engineering).

From an institutional point of view, theoretical engineering is an academic endeavour pursued by researchers (in aeronautical engineering, civil engineering, etc.) in the engineering departments

---

<sup>61</sup>Applied natural science can also investigate natural objects (e.g. forests in forestry science, or humans in medical science). Applied natural science can be considered both: as an *application* of science (since it *uses* results from theoretical natural science), and as applied *science* proper (since it *produces* knowledge for practical use), see section 2.1. As we have seen, Hansson does not distinguish those two senses.

of universities, technological institutes, research centres, etc. By contrast, practical engineering (performed by what we usually call ‘engineers’) typically belongs to the private sphere (industry and services). Of course there is not strict dichotomy: academic departments can have (and indeed do, as illustrated for example by research programmes such as the European Horizon 2020 framework) collaborations with private companies; and research can take place in private companies (in the form of research and development), whose employees can occasionally publish in academic journals. However, in many respects practical engineering resembles more an art or practice than a science – it can be considered an art which has been ‘scientificated’ (Niiniluoto, 1993, p. 11): engineers use scientific results, on-the-shelf software and/or in-house code which they have developed, rules of thumb and their own experience (which can be in the form of either tacit knowledge, practical rule knowledge (see Hansson, 2013), or even formalised in quantitative formulas), to design objects and solve practical problems. They do not, as a rule, create new scientific knowledge.

### 3.4.3 Examples

Let us now take Hansson’s (2018, p. 76) example of a ski lift whose cable is made of a new stainless steel alloy whose resistance to wear is higher than usual cables. This resistance has been scientifically established through prestigious peer-reviewed publications. The ski lift owner therefore asks for fewer security inspections of the cable than with previous cables. The question asked by Hansson is: should the security inspection follow the ski lift owner’s request? Hansson does not explicitly answer this question (presumably his answer would be affirmative), but instead claims that such a situation should not even happen, because of the continuous adjustment of the scientific corpus described above. It is unclear, however, in which sense this situation should not happen, i.e. how exactly the scientists and theoretical engineers working on this new alloy could and should anticipate its practical applications. It is also unclear how the practical engineers who design the ski lift would use this new knowledge. Finally, it is not certain that the inspection authority would indeed grant the ski lift owner’s request.

Let us take another example, in aeronautical engineering. This is a field which requires high safety levels, where numerical simulations and empirical tests (e.g. bench tests, wind tunnel and flight tests) prior to the certification of new aeronautical systems or designs are mandatory and represent an essential part of the professional practice. This is because the processes studied are usually too complex to allow for a mathematical solution (Hansson, 2013, p. 23)<sup>62</sup>, and because simulations and laboratory experiments are in turn never able to fully reproduce real-world conditions, with all their contingency. It seems that no aeronautical engineer would ever accept to skip these tests, because of the high stakes (human lives) and the value of safety which is paramount in aeronautics. On the contrary, aeronautical engineers have a tendency in their practice to add safety margins and/or factors (on this subject, see Möller and Hansson, 2008) to those already required by the regulation authorities. Even if practical engineering does not belong to science (where the corpus is produced) but to its application (where it is used), such professional practice can make us interrogate the relevance of the corpus model, putting into question the idea that the corpus is apt to any application. How exactly do such tests relate to the scientific corpus? Here one must distinguish:

- the available scientific knowledge (stemming from theoretical engineering and other sciences) used in practical engineering: for example, the handbook of hydraulic resistance (Idel’čik,

---

<sup>62</sup>Hansson writes about technological science (i.e. theoretical engineering), whereas my concern here is practical engineering. However, he also mentions automobile crash tests and endurance tests of furniture and household appliances, hence he seems to combine both fields.

1960/1966) used for the design of environmental control systems (ECS) and wing ice protection systems (WIPS) in aeronautical engineering;

- and the ‘knowledge’ or results produced by practical engineering itself, which concern *different, new* systems or designs (for example, a new environmental control system, or a new wing profile).

*Prima facie*, the former knowledge, which indeed belongs to the corpus, is relied on and not put into doubt in practical engineering; for example ‘the’ Idel’čik (1960/1966) represents the ‘bible’ of any thermal systems engineer. Only when *new* systems or designs are created (for example a new ECS or a new wing profile) will tests be carried out. This may lead to test again components (such as pipes, whose hydraulic resistance is provided by the Idel’čik (1960/1966)) whose individual behaviour is known in the literature, but such components will usually not be tested individually (precisely because of this reason), only as parts of a system (whose global behaviour has only been simulated and must be empirically checked). Similarly, the design of a new wing profile will rely on previously accepted scientific knowledge (such as Bernoulli’s equation in fluid dynamics), and only the new profile will be tested. The corpus model therefore seems applicable.

Nevertheless, the use of (multiplicative) safety factors and/or (additive) safety margins in order to create a safety reserve, i.e. a margin between the actual conditions of operation and the conditions which would lead to failure, is ubiquitous in engineering (e.g. Doorn and Hansson, 2017). In aeronautical engineering, it is common practice to over-dimension the structure as well as the systems of the aircraft. For example, the cooling capacity of the ECS, or the de-icing capacity of the WIPS, are dimensioned over the worst case scenarios. In civil engineering, safety factors are intended to compensate for five types of failure (e.g. Doorn and Hansson, 2017, p. 95):

1. higher loads than those foreseen,
2. worse properties of the material than foreseen,
3. imperfect theory of the failure mechanism in question,
4. possibly unknown failure mechanisms, and
5. human error (e.g., in design).

The first two types of failures are due to aleatory uncertainties (related to the non-deterministic nature of the world) whereas the last three are due to epistemic uncertainties (related to our incomplete or fallible knowledge). The failures #3 and (to a lesser extent) #4, which can be generalised to other engineering fields, are of relevance here, since they illustrate how engineers make use of scientific knowledge<sup>63</sup>, and the trust which they put in it. The fact that engineers use safety factors and/or margins to account for imperfect or incomplete scientific knowledge primarily shows that they do not consider this knowledge as fully reliable for their practical applications, otherwise no safety factor / margin would be necessary (for these types #3 and #4 of failure). In this sense, it could also be seen as a reason to increase the epistemic requirement of the claim(s) in question. On the other hand, the very fact that these safety factors / margins are applied on the basis of existing scientific knowledge shows that engineers nevertheless trust this knowledge to some extent, otherwise it would not serve as a basis of application of these safety factors / margins in the first place.

---

<sup>63</sup>Of course, engineering does not only rely on scientific knowledge, but also produces a knowledge of its own, as we have seen above.

In sum, it is not clear whether these two examples illustrate or not the corpus model. Even if the latter is more or less normative and that its non-application is not sufficient to invalidate it, it is also supposed to be an (idealised) description of scientific and engineering practice. Therefore, if it appears not to correspond to this actual practice, appears empirically hard to apply or has too many drawbacks in certain areas of science or engineering, this could be an argument for modifying or abandoning it. One of the main features of the corpus model is indeed that it is (supposed to be) universal, i.e. to apply to all areas of science and technology. More detailed empirical studies would certainly be helpful in order to better assess the actual relevance of the model.

## 4 Conclusion

In this article I have recalled the origins of Hansson’s corpus model of value influence; presented the model in detail, with its central feature of taking the maximum level of evidence required by the applications of a claim in order for the latter to enter the scientific corpus; presented the limits of the model; shown that it handles well controversial non-epistemic values; and concluded that it is a very good candidate for managing non-epistemic value influence, because contrary to other models in the literature it is fairly simple, it provides a universal rule for dealing with values, and it systematically preserves the epistemic integrity of science. I have then commented on some potential difficulties associated with the model’s central feature of taking the maximum requirement for corpus entry. Firstly, this maximum requirement can be difficult to identify. Secondly, the fact that this maximum requirement is non-optimal with respect to other less demanding requirements is well handled by the model, but it does not completely rule out the relevance of multiple corpora, depending on how the boundaries of the corpus are defined. Thirdly, the model depends on a presupposition of systematic preference for false negatives over false positives within science, which can have detrimental consequences (both within science and outside). Lastly, it is not entirely clear if there can be cases which require higher levels of evidence than this maximum requirement, thereby challenging the model. I have therefore called for further empirical work to assess this latter point.

Indeed, to reason abstractly about a model of value influence, without concrete examples, is difficult (even if such difficulty can be inherent in the work of abstract conceptual theorising), but most importantly it can also become pointless at some stage. Indeed, if we want the corpus model to inform not only the philosophical debate, but also scientists’ and engineers’ practice, and serve as a potential baseline for professional guidelines, we must ensure that it is realistic and can be applied. Even if the corpus model is normative, which means that its non-realisation should theoretically not be a reason for abandoning it, in practice if it appears too far from scientific or engineering practice, unrealistic or unfeasible, it can be a reason for modifying or abandoning it<sup>64</sup>. This is particularly challenging in the case of the corpus model which purports to be universal, regardless of the disciplinary field or practical application.

Therefore, improving the corpus model may benefit from empirical work, which may take the form of case studies, surveys or interviews, in order to help illustrate or revise the model, or inspire a new one (on the role of empirical work for philosophy of science, see Wagenknecht, Nersessian

---

<sup>64</sup>Some authors go as far as to say that scientific practice *justifies* rules of hypothesis acceptance (Hempel, 1981, §7): we should not look for a normative, ‘transcendent basis for inductive rules’; rather, ‘their grounding comes in part from their conformance to the practices and presystematic [i.e. intuitive] judgments [of scientists] they serve to codify and in part from the systematic advantages [the] codification [of such rules] provides.’ (Feleppa, 1981, p. 416)

and Andersen, 2015, p. 7; Steel, Gonnerman and O'Rourke, 2017). It seems reasonable, in order to have a normative stance on something, to first know what that thing concretely is. In order to come up with an operational model, scientists' and engineers' own normative views should especially be investigated, in order to explicitly distinguish between cases when scientists allow, or do not allow, non-epistemic values to have an influence (not just case studies showing that scientists are much more influenced by values than what they presumably believe). For example the views of researchers and engineers in aeronautical engineering about the relationship between scientific knowledge, consensus and chosen safety levels, may provide interesting material for potentially improving the corpus model. Case studies investigating scientists' own normative views are still very rare in the philosophy of science literature (see e.g. Gundersen, 2020)<sup>65</sup>, not to speak of engineers'.

## Acknowledgments

This work was supported by the Jacques Loeb Centre for the History and Philosophy of the Life Sciences (Ben-Gurion University of the Negev, Israel), the MTA Lendület Values and Science Research Group, and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement #101066502 (project SYSVAL). I thank Matti Eklund, Sven Ove Hansson, and the two anonymous reviewers (especially one of them for his very detailed review) for very helpful comments on this paper.

## References

- Anderson, Elizabeth (2004). "Uses of value judgments in science: A general argument, with lessons from a case study of feminist research on divorce". In: *Hypatia* 19.1, pp. 1–24. DOI: <https://doi.org/10.1111/j.1527-2001.2004.tb01266.x>.
- Beauchamp, Tom L. and James F. Childress (1994/2012). *Principles of biomedical ethics*. Oxford: Oxford University Press.
- Bridgman, Percy Williams (1947). "Scientists and social responsibility". In: *The Scientific Monthly* 65.2, pp. 148–154. URL: <https://www.jstor.org/stable/19013>.
- Bryse, Keynyn et al. (2013). "Climate change prediction: Erring on the side of least drama?" In: *Global Environmental Change* 23.1, pp. 327–337. ISSN: 0959-3780. DOI: <https://doi.org/10.1016/j.gloenvcha.2012.10.008>. URL: <https://www.sciencedirect.com/science/article/pii/S0959378012001215>.
- Bueter, Anke (2021). "Rationality and the Value-Freedom of Science". In: *Handbook of Rationality*. Ed. by W. Spohn { M. Knauff, giveni=W. S. bibinitperiod M. MIT Press, pp. 757–766.

---

<sup>65</sup>Gundersen's interviews of climate scientists lead him to recommend a case-by-case framework for articulating the VFI with other normative commitments, i.e. exactly the contrary of the approach promoted here: 'Instead of seeking to formulate one single general and categorical principle or distinction that prescribes what role non-epistemic values must play in the reasoning of scientific experts, one might consider whether a more promising approach is to lay out a set of principles that scientific experts must interpret and apply on a case-by-case basis in light of the other standards that they face.' (114) But of course this need not be the case. Common features between different interviewees could also be found, which would illustrate a universal model.

- Carrier, Martin (2022). “What Does Good Science-Based Advice to Politics Look Like?” In: *Journal for General Philosophy of Science* 53.1, pp. 5–21. DOI: <https://doi.org/10.1007/s10838-021-09574-2>.
- ChoGlueck, Christopher (2018). “The error is in the gap: Synthesizing accounts for societal values in science”. In: *Philosophy of Science* 85.4, pp. 704–725. DOI: <https://doi.org/10.1086/699191>.
- Churchman, C West (1948). “Statistics, pragmatics, induction”. In: *Philosophy of Science* 15.3, pp. 249–268. DOI: <https://doi.org/10.1086/286991>.
- Doorn, Neelke and Sven Ove Hansson (2017). “Factors and Margins of Safety”. In: *Handbook of Safety Principles*. Ed. by Niklas Möller et al. Vol. 9. Wiley Online Library, pp. 87–114.
- Douglas, Heather (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Elliott, Kevin (2011). *Is a little pollution good for you?: Incorporating societal values in environmental research*. OUP USA. DOI: <https://doi.org/10.1093/acprof:oso/9780199755622.001.0001>.
- (2017). *A tapestry of values: An introduction to values in science*. Oxford University Press.
- Elliott, Kevin and Daniel Steel (2017). *Current controversies in values and science*. Taylor & Francis.
- Elliott, Kevin C (2006). “An ethics of expertise based on informed consent”. In: *Science and Engineering Ethics* 12.4, pp. 637–661. DOI: <https://doi.org/10.1007/s11948-006-0062-3>.
- Elliott, Kevin C et al. (2017). “Values in environmental research: Citizens’ views of scientists who acknowledge values”. In: *PloS one* 12.10, e0186049. DOI: <https://doi.org/10.1371/journal.pone.0186049>.
- Elliott, Kevin C. (2013). “Douglas on values: From indirect roles to multiple goals”. In: *Studies in History and Philosophy of Science Part A* 44.3, pp. 375–383. ISSN: 0039-3681. DOI: <https://doi.org/10.1016/j.shpsa.2013.06.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0039368113000423>.
- (2022). *Values in Science*. Elements in the Philosophy of Science. Cambridge University Press. DOI: 10.1017/9781009052597.
- Feleppa, Robert (1981). “Epistemic utility and theory acceptance: Comments on Hempel”. In: *Synthese*, pp. 413–420. DOI: <https://doi.org/10.1007/BF01130052>. URL: <http://www.jstor.org/stable/20115603>.
- Guay, A. (2018). *Risque et expertise. 6es conférences Pierre Duhem*. Presses universitaires de Franche-Comté.
- Gundersen, Torbjørn (2018). “Scientists as experts: A distinct role?” In: *Studies in History and Philosophy of Science Part A* 69, pp. 52–59. DOI: <https://doi.org/10.1016/j.shpsa.2018.02.006>.
- (2020). “Value-free yet policy-relevant? The normative views of climate scientists and their bearing on philosophy”. In: *Perspectives on Science* 28.1, pp. 89–118. DOI: [https://doi.org/10.1162/posc\\_a\\_00334](https://doi.org/10.1162/posc_a_00334).
- Hansson, Sven Ove (2007a). “Values in pure and applied science”. In: *Foundations of Science* 12.3, pp. 257–268. DOI: 10.1007/s10699-007-9107-6.
- (2007b). “What is technological science?” In: *Studies in History and Philosophy of Science Part A* 38.3, pp. 523–527. DOI: 10.1016/j.shpsa.2007.06.003.
- (2008a). “Regulating BFRs—from science to policy”. In: *Chemosphere* 73.2, pp. 144–147. DOI: 10.1016/j.chemosphere.2008.01.073.
- (2008b). “Response to questions on hormesis and ethics”. In: *BELLE (Biological Effects of Low Level Exposures) Newsletter* 14.3, pp. 7–10.



- Hansson, Sven Ove (2010). “Changing the scientific corpus”. In: *Belief revision meets philosophy of science*. Ed. by Olsson Erik and Enqvist Sebastian. Springer, pp. 43–58. DOI: 10.1007/978-90-481-9609-8.
- (2013). “What is technological knowledge?” In: *Technology teachers as researchers. Philosophical and Empirical Technology Education Studies in the Swedish TUFF Research School*. Ed. by Inga-Britt Skogh and Marc J. De Vries. Brill Sense, pp. 15–31.
- (2014). “How Context Dependent Is Scientific Knowledge?” In: *Epistemology, Context, and Formalism*. Springer, pp. 127–140. DOI: 10.1007/978-3-319-02943-6\_8.
- (2017). “How values can influence science without threatening its integrity”. In: *Logic, methodology and philosophy of science—Proceedings of the 15th International Congress*, pp. 207–221.
- (2018a). “Politique du risque et intégrité de la science. 6es conférences Pierre Duhem”. In: *Risque et expertise*. Ed. by Alexandre Guay. Presses universitaires de Franche-Comté, pp. 57–86.
- (2018b). “Réponses. 6es conférences Pierre Duhem”. In: *Risque et expertise*. Ed. by Alexandre Guay. Presses universitaires de Franche-Comté, pp. 101–103.
- (2020a). “Social constructionism and climate science denial”. In: *European Journal for Philosophy of Science* 10.3, pp. 1–27. DOI: <https://doi.org/10.1007/s13194-020-00305-w>.
- (2020b). “Values in pharmacology”. In: *Uncertainty in Pharmacology*. Ed. by Adam LaCaze and Barbara Osimani. Springer, pp. 375–396. DOI: [https://doi.org/10.1007/978-3-030-29179-2\\_16](https://doi.org/10.1007/978-3-030-29179-2_16).
- Harsanyi, John C (1983). “Bayesian decision theory, subjective and objective probabilities, and acceptance of empirical hypotheses”. In: *Synthese* 57.3, pp. 341–365. DOI: <https://doi.org/10.1007/BF01064702>.
- Hempel, Carl G (1960). “Inductive inconsistencies”. In: *Synthese* 12, pp. 439–469. DOI: [https://doi.org/10.1007/978-94-017-2111-0\\_10](https://doi.org/10.1007/978-94-017-2111-0_10).
- (1981). “Turns in the evolution of the problem of induction”. In: *Synthese* 46, pp. 389–404. DOI: <https://doi.org/10.1007/BF01130050>. URL: <https://www.jstor.org/stable/20115601>.
- Hicks, Daniel J (2014). “A new direction for science and values”. In: *Synthese* 191.14, pp. 3271–3295. DOI: 10.1007/s11229-014-0447-9. URL: <https://www.jstor.org/stable/24026188>.
- Holman, Bennett and Torsten Wilholt (2022). “The new demarcation problem”. In: *Studies in history and philosophy of science* 91, pp. 211–220. DOI: <https://doi.org/10.1016/j.shpsa.2021.11.011>.
- Idel’čik, Isaak Evseevič (1960/1966). *Handbook of hydraulic resistance: coefficients of local resistance and of friction*. Israel Program for Scientific Translations.
- Intemann, Kristen (2015). “Distinguishing between legitimate and illegitimate values in climate modeling”. In: *European Journal for Philosophy of Science* 5.2, pp. 217–232. DOI: 10.1007/s13194-014-0105-6.
- Jeffrey, Richard C (1956). “Valuation and acceptance of scientific hypotheses”. In: *Philosophy of Science* 23.3, pp. 237–246. DOI: 10.1086/287489. URL: <https://www.jstor.org/stable/185418>.
- John, Stephen (2012). “Mind the gap. Review of Heather Douglas: Science, Policy and the Value-free Ideal”. In: *Studies in History and Philosophy of Science Part A* 43.1, pp. 218–220. ISSN: 0039-3681. DOI: <https://doi.org/10.1016/j.shpsa.2011.10.007>. URL: <http://www.sciencedirect.com/science/article/pii/S0039368111000884>.
- Kitcher, Philip (2011). “Philosophy inside out”. In: *Metaphilosophy* 42.3, pp. 248–260. DOI: <https://doi.org/10.1111/j.1467-9973.2011.01684.x>.
- Kourany, Janet A. (2010). *Philosophy of Science after Feminism*. Oxford University Press. DOI: <https://doi.org/10.1093/acprof:oso/9780199732623.001.0001>.

- Kuhn, Thomas S (1977). “Objectivity, value judgement, and theory choices”. In: *The Essential Tension—Selected Studies in Scientific Tradition and Change*. Chicago and London: The University of Chicago Press.
- Lacey, Hugh (2017). “Distinguishing between cognitive and social values”. In: *Current controversies in values and science*. Ed. by Kevin C Elliott and Daniel Steel. Routledge, pp. 15–30.
- Levi, Isaac (1960). “Must the scientist make value judgments?” In: *The Journal of Philosophy* 57.11, pp. 345–357. DOI: <https://doi.org/10.2307/2023504>.
- Lewandowsky, Stephan et al. (2015). “Seepage: Climate change denial and its effect on the scientific community”. In: *Global Environmental Change* 33, pp. 1–13. ISSN: 0959-3780. DOI: <https://doi.org/10.1016/j.gloenvcha.2015.02.013>. URL: <https://www.sciencedirect.com/science/article/pii/S0959378015000515>.
- Longino, Helen (1990). *Science as Social Knowledge. Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press. DOI: <https://doi.org/10.1515/9780691209753>.
- McMullin, Ernan (1982). “Values in science”. In: *PSA: Proceedings of the biennial meeting of the philosophy of science association*. Vol. 1982. 2. Philosophy of Science Association, pp. 3–28. DOI: 10.1086/psaprocbienmeetp.1982.2.192409. URL: <http://www.jstor.org/stable/192409?origin=JSTOR-pdf>.
- Merton, Robert K. (July 1938). “Science and the Social Order”. In: *Philosophy of Science* 5.3, pp. 321–337. DOI: 10.1086/286513. URL: <http://www.jstor.com/stable/184838>.
- Möller, Niklas and Sven Ove Hansson (2008). “Principles of engineering safety: Risk and uncertainty reduction”. In: *Reliability Engineering & System Safety* 93.6, pp. 798–805. DOI: [dx.doi.org/10.1016/j.ress.2007.03.031](https://doi.org/10.1016/j.ress.2007.03.031).
- Morgan, Gregory J (2010). “Heather Douglas: Is Science Value-Free?(Science, Policy, and the Value-Free Ideal)”. In: *Science and Engineering Ethics* 16.2, pp. 423–426. DOI: 10.1007/s11948-010-9204-8.
- Niiniluoto, Ilkka (1993). “The aim and structure of applied research”. In: *Erkenntnis* 38.1, pp. 1–21. DOI: <https://doi.org/10.1007/BF01129020>.
- Parascandola, Mark (2010). “Epistemic risk: empirical science and the fear of being wrong”. In: *Law, Probability & Risk* 9.3-4, pp. 201–214. DOI: 10.1093/lpr/mgq005.
- Polanyi, Michael (1962/2000). “The Republic of Science: Its Political and Economic Theory”. In: *Minerva* 1.38, pp. 1–21. DOI: <https://doi.org/10.1023/A:1026591624255>. URL: <http://www.jstor.org/stable/41821153>.
- Reiss, Julian and Jan Sprenger (2020). “Scientific Objectivity”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2020. Metaphysics Research Lab, Stanford University. URL: <https://plato.stanford.edu/archives/win2020/entries/scientific-objectivity/>.
- Rudner, Richard (1953). “The scientist qua scientist makes value judgments”. In: *Philosophy of science* 20.1, pp. 1–6. DOI: 10.1086/287231. URL: <https://www.jstor.org/stable/185617>.
- Ruphy, Stephanie (2006). ““Empiricism all the way down”: a defense of the value-neutrality of science in response to Helen Longino’s contextual empiricism”. In: *Perspectives on Science* 14.2, pp. 189–214. DOI: <https://doi.org/10.1162/posc.2006.14.2.189>.
- Steel, Daniel, Chad Gonnerman and Michael O’Rourke (2017). “Scientists’ attitudes on science and values: Case studies and survey methods in philosophy of science”. In: *Studies in History and Philosophy of Science Part A* 63, pp. 22–30. DOI: <http://dx.doi.org/10.1016/j.shpsa.2017.04.002>.

- UN (Mar. 1992). *Annex 1. RIO DECLARATION ON ENVIRONMENT AND DEVELOPMENT*. Rio de Janeiro. URL: [https://www.un.org/en/development/desa/population/migration/generalassembly/docs/globalcompact/A\\_CONF.151\\_26\\_Vol.I\\_Declaration.pdf](https://www.un.org/en/development/desa/population/migration/generalassembly/docs/globalcompact/A_CONF.151_26_Vol.I_Declaration.pdf).
- Wagenknecht, Susann, Nancy J Nersessian and Hanne Andersen (2015). “Empirical philosophy of science: Introducing qualitative methods into philosophy of science”. In: *Empirical Philosophy of Science*. Springer, pp. 1–10. DOI: 10.1007/978-3-319-18600-9.