REVUE
DE LA SOCIÉTÉ
DE PHILOSOPHIE
DES SCIENCES

Vol 12 N°1 2025 *https://doi.org/10.20416/LSRSPS.V12I1.3*

Anouk Barberousse & Marion Vorms

What is good scientific expertise for decision-making? Relevance matters

SOCIÉTÉ DE PHILOSOPHIE DES SCIENCES (SPS)
École normale supérieure
45, rue d'Ulm
75005 Paris
www.sps-philoscience.org

REVUE
DE LA SOCIÉTÉ
DE PHILOSOPHIE
DES SCIENCES

Anouk Barberousse & Marion Vorms

What is good scientific expertise for decision-making? Relevance matters

Sommaire

Cet article vise à clarifier les conditions du succès – ou les raisons de l'échec – de la communication entre experts scientifiques et décideurs, en mettant en avant l'importance de la pertinence des informations transmises pour les questions que se posent ces derniers. Il se concentre en particulier sur le cadre du tribunal, qui permet d'examiner à la loupe le rôle crucial que joue la pertinence des questions posées aux experts et des réponses qu'ils y apportent.

This article aims to clarify the conditions under which communication between scientific experts and decision-makers succeeds – or fails – by emphasising the importance of the relevance of the information provided in relation to the questions decision-makers are asking. It focuses in particular on the courtroom setting, which offers a close-up view of the crucial role played by the relevance of the questions addressed to experts and the answers they provide.

Mots clés : Expertise scientifique, Epistémologie Sociale, Décision, Pertinence, Témoignage.
Keywords: Scientific expertise, Decision-making, Relevance, Testimony

1. Introduction

Although its results are sometimes uncertain and subject to revision, science is widely regarded as the most reliable and objective source of information about the world. As a result, there is little doubt that a decision is more likely to be sound if it takes scientific knowledge into account rather than ignoring it. This is why decision-makers often rely on the epistemic authority of specialists – individuals or panels of experts – whom they task with answering a number of questions.

The use of scientific expertise to support decision-making raises several challenges, some of which have been clearly identified by philosophers of science. The first concerns the identification of 'good' experts: how can the credibility of a scientist be assessed, and how can one adjudicate between opposing views without being a specialist in the field oneself (Goldman 2001)? Who should we believe, and on what basis? On this point, it is widely acknowledged that non-specialists must rely on external markers such as institutional recognition, which serve as the most accessible indicators of credibility. However, unlike certain sociological approaches that reduce expertise to social markers (see Collins and Evans 2009), philosophers of science have focused on the relationship between these markers (and, more broadly, the institutional organisation of science) and the actual epistemic reliability of both scientific results and individual expert statements (Blais 1987; Hardwig 1991; Bouchard 2016).

Beyond the issue of the reliability of information provided by experts in a given field, there lies another question that is just as crucial for decision-making: that of usefulness. Expert reports are rarely sought for purely contemplative purposes – they are tools intended to inform decisions. What

REVUE
DE LA SOCIÉTÉ
DE PHILOSOPHIE
DES SCIENCES

requirements must a report meet to constitute good expertise, that is, to provide decision-makers with the information they need to make the best possible decision? Framed from the perspective not of the recipients (and principals) of expert advice but of the scientists called upon to serve as experts[1], this question raises the issue of the norms that govern expert conduct and the best practices they must adopt. How can they meet the demands of efficiency and usefulness expected of their mission? How can they provide decision-makers with information that will genuinely help them reach better decisions, given their specific goals? And since the knowledge produced by scientific research is rarely clear-cut or definitive, how can experts offer answers that are sufficiently clear and unambiguous to be understood and used, without sacrificing accuracy – that is, while remaining faithful to the current state of scientific knowledge?

This article aims to clarify the conditions for the success – or the reasons for the failure – of communication between experts and decision-makers, understood as a process aimed at putting the latter in the most favourable epistemic position to make the best decision in light of their own preferences (rather than those of the experts). We argue that this success is largely measured by the conclusions decision-makers draw from the statements provided by scientists. In other words, success depends not only on the accuracy of the information received but also on its relevance: what it enables decision-makers to conclude with respect to the specific issue at hand is just as important as whether it is correct.

From the decision-maker's perspective, then, the key question when interpreting an expert report is not only *Can I believe this?* but also *What can I infer from it for the decision I have to make?* From the expert's point of view, if they wish to be genuinely useful[2], one of the main questions they must ask is: how can I provide information that is relevant to the decision-maker, given the motivations behind the request for expertise and the role this information will play in the reasoning process that leads to a decision? This is not always easy to determine. Whether the issue is reducing greenhouse gas emissions, containing the spread of a virus, or continuing the exploitation of deep-sea resources, decision-makers must usually rely not on a single statement, but on a body of information – often drawn from different disciplines, marked by varying degrees of uncertainty, and sometimes pointing in conflicting directions. They must weigh and combine these elements to arrive at conclusions stable and robust enough to support their decisions.

In this respect, their position is, to some extent, comparable to that of judges or jurors in a criminal trial, who must evaluate a heterogeneous set of more or less credible – and sometimes contradictory – pieces of evidence, assess the probative value of each, and integrate them into the most accurate possible representation of events, in order to reach a fair decision. The approach we adopt in this article thus assumes that the tasks of judges and decision-makers are similar in kind: they must conduct complex epistemic enquiries (in a sense we will clarify), consisting of drawing practical conclusions from a body of heterogeneous and sometimes uncertain information produced by various actors. This kind of enquiry – what we might call 'evidential reasoning under uncertainty' – deserves to be studied in general, whether or not the information in question is scientific. Nevertheless, the specific challenges of such reasoning are intensified when scientific information is involved. The central aim of this article is to highlight these challenges and to draw out their implications for a reflection on best practices in scientific expertise.

Because our approach treats communication between experts and decision-makers as embedded within complex epistemic enquiries, we use testimony in criminal trials as our guiding model. The courtroom witness – whether expert or not – represents, in some respects, an oversimplified version of how expertise functions in public decision-making: it is typically individual scientists who testify in court, whereas reports that inform public policy are generally produced by panels of experts; moreover, the procedural constraints of judicial contexts often hinder, if not preclude, a healthy dialogue aimed sincerely and constructively at establishing the facts. Nevertheless, examining scientific testimony in court provides a magnifying lens through which to observe the various ways in which misunderstandings about relevance can arise in complex epistemic enquiries.

In the first part of the article, we define the notion of epistemic relevance and show – by focusing on ordinary (non-scientific) testimony in court – that a misunderstanding regarding the reasons a witness makes certain claims (i.e. the relevance they attribute to them) can undermine communication. We then examine how such misunderstandings may arise in the case of scientific testimony in court, and draw from this analysis a set of implications for the responsibilities of scientists who are called upon to act as experts.

1 It is important to distinguish between two meanings of the term 'expert'. In the first sense, it refers to a specialist in a scientific field. In the second, it refers to the role such an individual (an expert in the first sense) may be called upon to play when non-specialists seek scientific insight (an expert report) to assist them in making a decision.

2 Throughout this article, we assume that experts are cooperative and act in good faith. The communication challenges we highlight between interlocutors acting in good faith have their counterparts in contexts involving intentional manipulation.

REVUE DE LA SOCIÉTÉ DE PHILOSOPHIE DES SCIENCES

2. Epistemic relevance

Relevance, as we understand it here, is a property of certain pieces of information[3] considered by an agent engaged in an epistemic enquiry – a property that depends both on the state of the world and on the agent's state of uncertainty or ignorance. By 'epistemic enquiry', we refer to the process of formulating and confirming one or more hypotheses on the basis of a set of non-conclusive factual information. Typical cases include criminal investigations, scientific research, and even more mundane activities such as trying to answer questions like *Where did the water that soaked the living room wall come from?* or *Why did the bathroom light suddenly go out?* A judge or juror evaluating the plausibility of the hypothesis that a defendant is guilty is likewise conducting an epistemic enquiry. A piece of information is relevant if it allows us to make progress in the enquiry – either by suggesting new hypotheses or by altering the plausibility of one of the hypotheses under consideration.

Understood in this way, the notion of relevance – which originates in common law reflections on evidence (Thayer 1969; Roberts 2022) – differs from the linguistic concept popularised by the work in pragmatics by Grice (1995) and then Sperber and Wilson (2001). In the linguistic framework, relevance is a property of speech acts, whereas the kind of relevance that concerns us here – epistemic relevance – applies to pieces of information or items of evidence, whether or not they are linguistic in nature, and regardless of their source. In section 1.2, we will see how these two forms of relevance come together in the context of courtroom testimony.

Some epistemic enquiries are particularly complex, in the sense that they involve a multitude of pieces of information whose connections – both with one another and with the main hypotheses under consideration – are neither unambiguous nor certain. It is in such contexts that assessing the epistemic relevance of the available information becomes both important and difficult. The fact-finding task of judges and jurors in a criminal trial – examining the pieces of a jigsaw puzzle in order to reach conclusions about a number of factual questions – is paradigmatic of complex epistemic enquiries. In this section, we use that model to explore the properties of epistemic relevance.

For the sake of simplicity, we use examples in this section involving non-scientific information. After defining the notion of relevance as it applies to any item of evidence that may play a role in evidential reasoning, whether testimonial or not (1.1), we will show how, in the case of testimony, the intentional dimension of communication adds a layer of difficulty in assessing relevance and, thereby, the probative value of the information conveyed (1.2).

2.1 The relevance of evidence in fact-finding

In what follows, the picture we present of the judge's or juror's fact-finding task in a criminal trial is highly abstract and idealised. It is not intended to correspond to actual situations, whether in common law systems or in continental legal systems. The figure of the judge or juror serves here primarily as a model for examining evidential reasoning and decision-making under uncertainty. Our analysis draws heavily on the work of David Schum (1994), who seeks to provide a general theory of evidential reasoning under uncertainty, focusing specifically on such reasoning in criminal investigations and trials.

2.1.1 Evidence, hypothesis, uncertainty, and decision

Let us consider the juror in a criminal trial, tasked with determining whether the accused is guilty. To do this, they rely on a set of evidence, based on which they must assess the plausibility of the main hypothesis (***H***), namely that the accused is indeed the perpetrator of the alleged acts.

H is irreducibly uncertain: even if the available evidence seems overwhelming, it is not and cannot be conclusive. Doubt is always, in principle, possible – we can imagine the existence of unknown elements that might weaken the evidence considered so far, and thus weaken ***H***; likewise, we can imagine an alternative hypothesis, incompatible with ***H***, that could also account for the known facts.

Despite this uncertainty, the jury is required to reach a categorical conclusion regarding ***H*** – that is, to ultimately accept or reject ***H*** in order to declare the accused guilty or not guilty. As with any decision (see, for example, Jeffrey 1996), the choice to accept or reject a hypothesis depends not only on epistemic parameters – namely, the assessment of ***H***'s plausibility given the available evidence – but also on non-epistemic parameters, including the stakes of the decision, particularly the consequences of the different types of possible error (e.g. convicting an innocent person or acquitting a guilty one). According to this model, the relative

3 Here, we use the term 'information' in its most common and neutral sense (distinct from, for example, the meaning theorised by Dretske 1981): that is, as an element that carries propositional content about the world. This element typically takes the form of a statement, but may also consist of a video image, a photograph, or a physical object, provided it is situated within a context that indicates what it is meant to tell us about the world. Our use of the term is neutral with respect to the truth value of the content in question (unlike Dretske's notion of information, which presupposes truth).

seriousness of each type of error helps determine a threshold of plausibility beyond which ***H*** should be accepted.[4]

In what follows, we will not focus on the question of that threshold, but rather on how the plausibility of a hypothesis is evaluated. In other words, our concern here is not the decision itself, but the reasoning that precedes it. Nevertheless, it is important to note that evidential reasoning – whether by a juror in court or by public authorities faced with high stakes, time-sensitive decisions under uncertainty – is always oriented toward a decision.[5]

2.1.2 A network of evidence

As crime novels and films often illustrate, the various pieces of evidence used to come up with and assess hypotheses about the perpetrator of a crime cannot be considered in isolation – they must be considered in relation to each other. This means that the probative value of each piece of evidence partly depends on the other available elements. Some pieces may seem directly incriminating or exonerating (e.g. a threatening letter, fingerprints at the crime scene, or conversely, proof that the suspect was thousands of kilometres away at the time of the murder) in the sense that they affect the plausibility of H rather straightforwardly, even on their own. Others do so only indirectly – for example, by bearing on the credibility of a witness or on how to interpret another, more directly relevant piece of evidence. For instance, learning that the accused visited the victim's home for an 'innocent' reason reduces the probative value of the discovery of their fingerprints there.

This way of linking a piece of evidence to ***H*** – either through a linear chain (which may contain several links) or by reassessing how another piece of evidence affects ***H***'s plausibility – is central to the kind of evidential reasoning that characterises complex epistemic enquiries. It is an essentially dynamic process, since the probative value of each element depends, as we have just seen, on how it connects with the others. Following David Schum (1994), who draws on categories from common law, we can analyse probative value along two key dimensions: *credibility* and *relevance*.

2.1.3 Credibility and relevance

Following David Schum, let us denote ***E**** a piece of evidence, obtained through a specific process and presented in a particular form – for example, a written DNA analysis report produced by a lab comparing a sample taken by a forensic doctor from under the victim's fingernails with DNA taken from the suspect by police during custody; a photograph of the victim taken by the police when the body was discovered; or a neighbour's account of an encounter with the suspect in the stairwell. ***E*** corresponds to the event or state of affairs represented by ***E****: the degree of DNA match, the position and location of the body, or the fact that the neighbour encountered the suspect in the stairwell. In other words, ***E*** is the propositional content of ***E****.

The credibility of ***E**** can be defined as the extent to which ***E**** indicates the truth of ***E*** – that is, the extent to which it warrants accepting ***E***. If ***E**** is highly credible, we may treat ***E*** as established and proceed with our reasoning accordingly. The credibility of a piece of evidence depends largely on its source[6], and the way it is assessed varies depending on the type of evidence. For a document or video, we check for authenticity; for DNA analysis, we assess the chain of custody and the precision and accuracy of the analysis; for testimony, we evaluate the witness's sincerity and competence (see Schum 1994, pp. 98–109).

The relevance of ***E**** lies in the way our knowledge of ***E*** helps us advance in evaluating ***H***. More precisely, ***E**** is relevant if the event or fact ***E*** it represents informs us about the plausibility of ***H***. This relevance can take different paths. A report showing a DNA match between what was found under the victim's fingernails and the suspect's DNA is clearly relevant to assessing guilt. A report about the weather in London on the day the Queen died, on the other hand, appears irrelevant – unless the consistency of a key witness's testimony depends on the weather that day; in that case, the

4 *Common law, with its standards of proof – particularly the standard of* beyond a reasonable doubt *– allows us to conceptualise the distinction between evaluating the plausibility of a hypothesis and making an epistemic decision to accept or reject it, a distinction that is absent from the French notion of* intime conviction.

5 *Even in a 'purely' epistemic enquiry – that is, one not intended to lead to any action in the world – epistemic decisions must be made, in the sense that, both throughout the enquiry and at the point of considering it complete, one must accept or reject hypotheses that remain irreducibly uncertain. Since Hempel (1960), the risk of error associated with accepting (or rejecting) a hypothesis about the world has been referred to as 'inductive risk'.*

6 *Another dimension of the acceptability of a propositional content* **E** *is its* a priori *plausibility (i.e. prior to the presentation of* **E***). A piece of information that appears highly implausible in light of our background beliefs or the other information available in a given case is unlikely to be readily accepted, even if it comes from a credible source. Moreover, the plausibility of a piece of information and the credibility of its source are not independent: implausible information gains plausibility when presented by a credible source, and conversely, low plausibility can diminish the perceived credibility of the source. See Bovens and Hartmann (2003) for a Bayesian model of how the content of an argument affects assessments of source credibility. For a psychological study of the interaction between these two dimensions, see Hahn, Harris, and Corner (2009). In this article, we set aside the question of the* a priori *plausibility of* **E**.

REVUE DE LA SOCIÉTÉ DE PHILOSOPHIE DES SCIENCES

weather report tells us something about the credibility of the testimony, and thus indirectly about the accused's guilt.[7]

The distinction between credibility and relevance applies to all types of evidence, whether testimonial or not. Unlike credibility, the relevance of a piece of evidence appears to be independent of both its source and its medium. Images from a surveillance camera and a security guard's account may differ in credibility, but if they represent the same thing (e.g. the suspect's presence at a doorway, dressed in a certain way, at a certain time), they share the same relevance. However, as we will see in section 1.2, assessing relevance – and distinguishing it from credibility – becomes more complex in the case of testimonial evidence, opening the door to various kinds of misunderstanding (or manipulation).

2.1.4 Properties of relevance

As defined above, the relevance of a piece of information depends both on the stage of the epistemic enquiry in which it is considered and on the other information available. Since relevance is measured by its capacity to change our assessment of ***H*** and of other possible hypotheses, it is relative to the framework of the epistemic enquiry – partially defined by ***H***. Information that is irrelevant within one framework may prove crucial in another. Likewise, the relevance of a piece of information may be significantly altered by the introduction of new elements or by the formulation of a novel hypothesis, which reorganises the network of inferences connecting the available elements. For instance, testimony placing a suspect at the cinema at the presumed time of a murder serves as an alibi, but may lose its relevance – or even acquire new relevance as incriminating evidence – if new information revises the estimated time of the crime and suggests that the perpetrator could have gone to the cinema before meeting the victim.

Moreover, to judge a piece of information as relevant is to say that it allows one to draw conclusions that modify the assessment of the plausibility of a hypothesis, whether it be ***H*** or one of the other hypotheses considered in the enquiry. Such conclusions are the result of inductive inferences – and are therefore uncertain – and they also rely on implicit assumptions, often in the form of empirical generalisations. These generalisations vary in how well founded and consciously held they are, and they may stem from common sense, the agent's personal experience, or scientific knowledge. For example, a medical report attesting to the accused's chronic alcoholism increases the plausibility of the hypothesis that he violently assaulted his partner, by drawing on the generalisation that alcohol consumption fosters violent behaviour.

For all the reasons mentioned above, the assessment of relevance depends on the agent's point of view – a point of view that may itself evolve over the course of the enquiry. Even assuming that all the jurors in a trial possess exactly the same information about the case, each of them, through personal experience, has acquired knowledge and formed representations of the world – that is, has generated empirical generalisations – that differ, at least in part, from those of the others. One juror, having come from the social environment in which the events took place, may perceive the relevance of certain details in a testimony that go unnoticed by others. Another, having conducted sociological research in the same milieu, may be more guarded against common but empirically unfounded representations.

Recognising this is neither to fall into relativism nor to abandon the idea of a standard of correct reasoning. Some inferences are logically flawed; others are based on empirical generalisations that are false or contradict other implicitly held assumptions. Still, evidential reasoning is inductive in nature, and multiple valid inferential paths are possible. One of the virtues of collegiality in a court is precisely that it enables individual experiences to be shared, implicit generalisations to be made explicit, and – through dialogue – a collective assessment to emerge regarding what it is most reasonable to conclude from the available information, based on its relevance and credibility.

2.2 The epistemic relevance of testimonial evidence

As announced above, we now turn to examine how the distinctions and proposals introduced thus far shed light on the task of judges and jurors faced with courtroom testimony. Most epistemological discussions of testimony – understood in a broad sense, far beyond the formal legal context – tend to focus on its credibility. These enquiries address the foundations of our trust in others as sources of knowledge; in other words, the principles under which it is rational to accept as true the content asserted by others (Burge 1993; Audi 1997; Fricker 1995; Coady 1995). When attention shifts to more practical concerns – whether to grant or suspend such fundamental trust in particular cases – these considerations are almost invariably framed in terms of the credibility of the witness or the plausibility of the content.

For instance, if Ms X says, *I passed the accused in the hallway*, the first question to arise is whether what she says is true (i.e. whether she really did encounter the accused there). That question, in turn, depends on her credibility, traditionally understood in terms of sincerity or good faith (the disposition to say what one believes to be true) and competence (the

7 We are here within the framework of a trial, where a well-defined hypothesis is being evaluated – namely, that the defendant is the author of the acts they are accused of. However, the notion of relevance can also apply to information that advances an enquiry at a stage where no specific hypothesis has yet been formulated. In such a context, relevant information is that which suggests hypotheses – that opens up lines of enquiry.

REVUE DE LA SOCIÉTÉ DE PHILOSOPHIE DES SCIENCES

capacity to form true beliefs). In all such cases, the central question regarding the epistemic value of testimony is: *Can I accept the content being reported to me?*

Yet the relevance of testimonial evidence raises distinctive problems that merit closer attention – and may even challenge the boundary between credibility and relevance.[8] Unlike other forms of evidence, whose relevance can be assessed independently of how they were obtained and regardless of their credibility, testimony is a communicative act. As such, it involves both the transmission of propositional content and the expression of communicative intent. As shown in the work of Grice (1995) and later Sperber and Wilson (2001), a core presupposition shared by participants in any communicative exchange is that all acts – linguistic or otherwise – are relevant to the meaning of the content exchanged. While pragmatic theories of relevance primarily aim to explain how the maxim of relevance guides the interpretation of utterances, the presumption of relevance inherent in any assertive act has further implications for the attribution of epistemic relevance by the agents engaged in enquiry. The information offered by a witness is interpreted by the hearers as being relevant to their epistemic concerns – if only because testimony in court is, by nature, a communicative act, and thus subject to the principle of relevance.[9]

In line with this principle, and despite the formulaic oath, witnesses on the stand are not expected to tell *the whole* truth (a notion that is in any case difficult to define), but rather *the relevant* truth. That is, they are expected to select information based on their assessment of what is relevant to the specific questions at stake in the enquiry to which they are contributing. This observation points to several difficulties that are specific to the transmission of information from witnesses to judges.

First, there is no guarantee that the witness and the judges or jurors will share the same assessment of what is relevant. Not only do they have different personal experiences, but they also possess different information about the case itself. This problem, which can be generalised beyond the legal context, becomes especially acute in trials, where witnesses are deliberately kept away from proceedings and ideally know as little as possible about the case – precisely to preserve their credibility. This isolation is intended to reduce bias (by preventing witnesses from selecting information that aligns with what they may wish to believe or make the judge believe) or to avoid their being influenced by external knowledge.[10] But given this disparity in information, how could witnesses possibly evaluate the relevance of their statements in the same way as the judges?

One might argue that this asymmetry is of little consequence. So long as the witness's statements are accurate, it is up to the judges – occupying an epistemically 'superior' position, since they know the broader context – to determine their relevance. In this view, as with other types of evidence, judges must ask whether the witness is credible, i.e. whether $\boldsymbol{E}^*$ licenses them to admit $\boldsymbol{E}$; it is then up to them to draw from $\boldsymbol{E}$ whatever they find useful. However, this approach treats the witness as a neutral or transparent source, and downplays a central lesson of pragmatics: that the reception of a message is inseparable from how the hearers interpret the speaker's communicative intentions. In conveying information, the witness not only asserts its accuracy but also implicitly suggests that it is relevant in the context of the trial – and further suggests that no other relevant information is being withheld. Upon receiving the testimony, the judge – despite having the authority to evaluate relevance independently – is likely to make inferences, perhaps unconsciously, about the relevance the witness intended to communicate. Such inferences may themselves be shaped by the judge's perception of the witness's personality, or by the position the witness appears to occupy in the proceedings – whether testifying for the prosecution or the defence. These inferences also lead the judge to attribute certain epistemic states to the witness – that is, to form hypotheses about what the witness knows, believes, or is unaware of, both regarding the case and more broadly.

Consider, for example, a witness who appears to be intent on establishing the accused's innocence (perhaps because they are a friend of the accused, or because they explicitly say they believe in the accused's innocence and present their testimony as proof). Suppose this witness claims to have seen the accused at a neighbourhood cinema during the 8 p.m.

8 *This is the subject of an article currently in preparation.*

9 *Moreover, the act of testifying, insofar as it is carried out by someone who may have a connection to the case, also constitutes an event within that case (it belongs both to the 'story of the crime' and the 'story of the trial', to borrow the terms used by Lagnado and Gerstenberg 2017, pp. 577–579). Not only do the witness's identity and their relationship to the accused and to the case inform our assessment of the credibility of their statements, but the very fact that the witness makes those statements (regardless of whether we deem them credible) is itself a potentially relevant event. Learning that a witness has lied, for example, may reveal more about the case than the content of the (discredited) testimony itself (see Lagnado and Harvey 2008 for a psychological study). Thus, as Christine astutely understood in Agatha Christie's* Witness for the Prosecution*, the discrediting of a prosecution testimony from an unfaithful wife may itself constitute (as an event) far more relevant – and exculpatory – evidence in the eyes of the jurors than the exculpatory testimony of a loving spouse. However, in this article, we focus on the relevance of the content of the witness's testimony, assuming the witness to be cooperative.*

10 *Conversely, in most legal systems, jurors are instructed to consider only the evidence presented to them, excluding any information from outside sources – particularly the media. Adhering to this requirement is, of course, much more difficult to enforce and monitor than restricting witnesses' access to the case file or the courtroom proceedings. Nevertheless, this further contributes to the asymmetry in the epistemic positions of witnesses and jurors with respect to the case.*

showing. The very fact that they report this event suggests that they consider it relevant. As a friend (and thus likely to testify in favour of the accused), we may infer that they would not spontaneously report an incriminating fact. It is therefore reasonable to assume that they see the accused's presence at that time and place as exculpatory – that is, as an alibi. Consequently, we may infer that they believe the crime took place that evening (unless we construct more elaborate hypotheses about the case and the witness's knowledge or beliefs).

But if this assumption turns out to be false, the detail provided may, unbeknownst to the witness, become incriminating. Still, in presenting this detail as an exculpatory one (or so it is presumed), the witness also offers indirect evidence about his broader epistemic position, which in turn frames the interpretation of the rest of his testimony. Yet if the judge misinterprets the witness's communicative intention – and thus makes mistaken assumptions about what the witness knows, believes, or ignores – these errors may lead to further epistemic mistakes, including factual errors about the case itself.

A simple example might be a judge interpreting a witness's silence about an apparently relevant event as evidence that the event did not occur – on the assumption that the witness would not have failed to mention it if it had. But perhaps the event *did* occur and the witness *is* aware of it, yet failed to see its relevance. If such a situation is coupled with a misinterpretation by the judge of the relevance the witness assigns to what they do report, genuine misunderstandings may arise, leading not only to errors about the witness's mental states but about the facts of the case themselves.

Despite such risks of misunderstanding – or, more broadly, communicative failure – the structure of the courtroom is such that the burden of determining the relevance of testimony seems to fall largely on the judges. In addition to their epistemic advantage, judges (and lawyers) are empowered to ask any clarifying questions they deem necessary. By contrast, witnesses are expected to provide answers without asking questions themselves.[11] Thus, if ambiguities or grey areas persist – unless deliberately maintained by the witness[12] – the responsibility for communication seems to lie with the legal professionals.[13] They are in a better position than the witnesses to assess the relevance of testimony; in this respect, they are epistemically superior, since they possess information to which the witnesses lack access, and they retain full freedom to request further clarification.

The situation differs, however, in the case of scientific testimony. Although the judge remains epistemically superior with regard to the details of the case – factual and legal alike – and retains control of the proceedings, judge and expert are not epistemic peers when it comes to the content of the testimony itself. Not only is the judge unable to access firsthand the content being reported (as with the lay witness, who testifies to events the judge did not witness), but – as Hardwig (1985) has shown – the judge is also not in a position to assess the methods used to validate and justify the scientific conclusions presented.[14]

What implications does this epistemic asymmetry have for the reception and use of scientific testimony? And, from the scientist's perspective, what responsibilities arise in choosing what information to communicate to the court?

3. Scientific testimony in court

Unlike the lay witness – who generally figures in the 'story of the crime' in a broad sense, insofar as they are connected to the events (even if only, say, by being the cinema ticket clerk where the accused claims to have been at the time of the crime) – the expert is not expected to be involved in the case

11 This asymmetry in the 'dialogue' between the questioners (judges and lawyers) and the witness prevents the latter from requesting clarification of the question posed – and, a fortiori, of the reasons behind it – even with the sincere aim of providing the most relevant answer. Beyond the fact that such a request would likely go unanswered, it could also damage the witness's credibility. It may be interpreted as an attempt to avoid or sidestep the question, or as a way of ensuring the answer aligns more closely with the interests the witness is defending rather than with the pursuit of truth (which the obligation to answer the question as it is posed *is meant to promote).*

12 As noted earlier, we assume good faith and a genuine cooperative effort on the part of the witness. Naturally, however, each possible path to misunderstanding also presents an opportunity for strategic manipulation.

13 Jennifer Saul (2012, p. 96), drawing on arguments by Solan and Tiersma (2005), defends the view that speakers bear less responsibility for how their statements are interpreted by listeners in the legal context than they do in ordinary life. Here, we extend that idea to include not only the interpretation of the content of the witness's statements, but also their epistemic relevance.

14 The layperson's epistemic dependence on the expert is characterised by the fact that the layperson '(1) has not performed the enquiry that would provide the evidence for his belief that p, *(2) is not competent, and perhaps could not even become competent, to perform that enquiry, (3) is not able to assess the merits of the evidence provided by [the] expert […]'s enquiry, and (4) may not even be able to understand the evidence and how it supports [the expert's] belief that* p' *(Hardwig, 1985, p. 339).*

in any way.[15] On the contrary, their complete detachment from the case is desirable, as any connection could jeopardise their neutrality. Experts are called to testify not because of their relation to the case itself, but in virtue of their privileged epistemic position with respect to the scientific or technical issues at stake.

This leads to another key difference from lay testimony: the expert's epistemic advantage includes a capacity to grasp the relevance of their statements – something judges do not necessarily possess. Scientists are better equipped than judges to determine what can legitimately be inferred from scientific data, since their work largely consists in examining the extent to which, and the conditions under which, various claims can be justified. Hardwig (1985) has shown that the epistemic dependence of non-experts (in this case, judges) on scientists compels them to trust experts not only with regard to justification (and thus the accuracy) of scientific claims, but – according to the account of epistemic relevance we propose here – with respect to their relevance as well. As a result, it is reasonable to hold that, unlike in the case of lay testimony, the responsibility for assessing the relevance of the content largely falls on the expert, rather than the judge. This assumption underlies the analysis that follows, in which we adopt the perspective of the scientific witness and ask how they ought to proceed in order to offer relevant responses.[16]

A clarification is in order regarding the types of scientific testimony in court that we are concerned with. Experts are frequently called upon to address specific questions – such as the match between two DNA traces, the trajectory of a bullet, or the compatibility between certain injuries and the hypothesis of poisoning. These forms of forensic expertise raise their own set of problems, whether due to misjudgments about the reliability of the investigative methods employed (e.g. graphology), or due to difficulties in the understanding and correct application of probabilities and statistics.[17] Our focus here, however, is on a different type of expert testimony: cases in which scientists are asked to provide assessments of general hypotheses – for instance, the toxicity of a particular molecule – in order to inform a legal decision. These situations, particularly common in public health trials, closely resemble certain forms of expertise intended to guide public policy decisions, which also often concern general scientific hypotheses.

After offering a preliminary analysis of the respective roles of credibility and relevance in expert testimony (section 2.1), we will use a simple fictional example to highlight several normative constraints that apply to scientific witnesses in delimiting the relevance of their statements for the court (section 2.2).

3.1 Credibility and relevance in scientific testimony

As with any form of testimony, the credibility of scientific testimony appears to involve two dimensions: good faith and competence. In principle, these are governed by the internal structure of scientific institutions, which are tasked with selecting competent scientists, ensuring the publication of sound work, and maintaining oversight of researchers' integrity and probity.

In the context of a trial – which is, by its nature, characterised by conflicting interests, and where the parties involved may pursue goals that do not always align with the search for truth – the issue of good faith takes on particular importance. It is obvious, for instance, that the defence is unlikely to summon a scientist whose conclusions are unfavourable to its case. Still, it would be a mistake to infer from this that the testimony of party-appointed scientists lacks any credibility. The judge must instead weigh the various markers of credibility – a task for which they may not always be well equipped. That said, we will set this issue aside for the moment, and our analysis of relevance will, from another angle, bring us back to questions of credibility and trust.

At first glance, one might assume that the problems we previously identified – where the witness and the judge diverge in their interpretation of what is relevant – are less likely to arise in the case of scientific testimony. Scientific

*15 In the French inquisitorial system – unlike common law adversarial systems, where all experts have the status of witnesses and are called by one of the parties – only court-appointed experts (*experts judiciaires*) hold the specific legal status of 'experts'. This status grants them access to the case file and permission to attend hearings prior to their testimony, and it requires them to take an oath that differs in nature from that of other witnesses. Scientists summoned by the parties are merely 'knowledgeable witnesses', like lay witnesses, they are excluded from the hearings and lack a comprehensive view of the available information – let alone the legal details of the case. We ignore these procedural details here, and will treat all scientists called to testify as witnesses – referring to them as 'experts'.*

16 We do not suggest that responsibility for successful communication lies entirely with scientists. Given the complexity of certain epistemic enquiries undertaken by decision-makers, the task of assessing the relevance of the questions posed to scientists – and of the answers scientists can provide – cannot be unilaterally assumed by scientists alone. Rather, it should ideally be the subject of an open dialogue (though such dialogue is difficult to establish in a courtroom setting).

17 A case in point is the error known as the 'prosecutor's fallacy', which involves inferring from $p(E|\neg H) = x$, *that* $p(H|E) = 1 - x$. *For example, drawing the conclusion that the probability of the defendant's guilt is 0.9 based given that the probability of the evidence under the assumption of innocence is 0.1. Once an event that is unlikely under the innocence hypothesis has occurred, this fallacy amounts to treating it as definitive proof of guilt. More broadly, errors stemming from the misuse or misinterpretation of conditional probabilities – such as neglecting base rates or overlooking the non-independence of certain probabilities – have been widely criticised. Notable cases include the trials of Lucia de Berk in the Netherlands and Sally Clark in the United Kingdom (see Schneps 2018). Some efforts, particularly in the UK, aim to clarify and standardise the use of Bayesian probabilities in forensic science (Cook et al. 1998; but see Fenton et al. 2014 for a caution against overly simplistic interpretations of the principles set out in this literature).*

questions may seem more closed-ended: a scientist asked about the toxicity of a molecule appears to face a more narrowly defined task than, say, a neighbour of a murder victim asked to report everything they may have observed that is (potentially) connected to the case. On this view, the primary virtue expected of scientific expert testimony is the credibility in the strict sense of its content – that is, its accuracy.

However, this view rests on a highly idealised and somewhat naïve conception of the kind of answer science can provide to decision-makers. When decision-makers call upon experts, their aim is not to acquire knowledge for its own sake, but to mobilise it in order to reach the best possible decision. The ideal scientific contribution would consist in a clear and intelligible presentation of all the information relevant to the decision at hand, in a form suited to the decision-maker's needs.

If such a body of information were readily available – if science could deliver categorical answers to every question decision-makers pose – then perhaps the only issue would be to assess the credibility of the experts offering these answers. But in most cases, clear-cut, definitive, and unambiguous answers are simply not available. Scientific findings are often marked by uncertainty, may lack consensus, and are frequently incomplete, leaving residual pockets of ignorance. There is, therefore, no unique way to answer a decision-maker's question that is both accurate and exhaustive. Like the lay witness, the expert cannot claim to tell 'the whole truth'. Rather, the expert must draw on the state of the art to select the information that comes closest to supporting the best possible decision – choosing relevant data and leaving aside what is not.

But how should the expert determine what is relevant to the decision? On what basis should they distinguish between what is and is not worth bringing to the court's attention? And what guidelines should they follow in presenting this selected information in a way that makes it intelligible and usable for non-specialist decision-makers?

As we shall see, the criteria that define a good scientific answer in the context of academic science are not always well suited to the expectations of a non-expert audience. Therefore, even assuming that the scientist is sincere, neutral, and competent, the challenges of selecting and presenting information make the task of scientific expertise in court a particularly delicate one.

3.2 Non-interference, exhaustiveness, and relevance

In this section, we begin with a highly simplified fictional example (2.2.1) to highlight certain requirements that fall upon the expert when they are called to provide input to support a decision. In particular, we will show that – contrary to a simplistic image of the expert's role, according to which they must avoid involvement in the extra-scientific inferences that their interlocutor may draw from the information provided – the expert must sometimes speak explicitly to the relevance of their input in order to assist the judge's reasoning (2.2.2). We then question the ideal of exhaustiveness, showing that the requirement of relevance can at times necessitate omitting certain information – specifically when such information risks misleading the judge by prompting illegitimate inferences (2.2.3).

3.2.1 A fictional example

Consider a fictional trial in which a pharmaceutical company is accused of marketing a drug compound allegedly involved in the deaths of several patients. Let us suppose that a key issue in determining the outcome of the trial is whether there exists a generic causal link[18] between the consumption of the compound and a particular type of cardiac injury.[19]

The defence calls on epidemiologists, who testify that the available data does not establish such a causal link. The prosecution, arguing that the compound is cardiotoxic, points to the weak evidential value of these findings, which come from retrospective cohort studies[20] marked by several methodological shortcomings. Seeking to convince the judges of the compound's toxicity, the prosecution also calls on toxicology experts, who report that laboratory experiments have robustly demonstrated severe cardiac effects in rats. These are the only conclusive experimental results. The defence counters by emphasising that metabolic differences between rats and humans preclude a straightforward extrapolation of these results to humans.

First, how should we assess the credibility of the scientific findings presented? The answer depends on their reliability – that is, the quality of the experiments underlying the analyses, the rigour and precision of the analyses

18 Generic causality, as opposed to particular causality, refers to the assumption of a causal link – usually probabilistic – between properties. 'Breathing asbestos dust causes cancer' is a generic causal claim; 'Mr. X's cancer was caused by asbestos' is a particular causal claim.

19 It should be emphasised that it is clearly not the role of a criminal judge to determine a question such as the toxicity of a given molecule. When such questions arise in a judicial context, it is generally in order to assess the responsibility of certain actors in the occurrence of a public health crisis. This typically involves examining issues related to their intentions and beliefs – that is, what they could or should have known or suspected regarding the toxicity in question.

20 According to the hierarchy established by evidence-based medicine (Sackett et al. 1996; Stegenga 2018), cohort studies are considered to yield less robust results than randomised clinical trials.

themselves, and the presence (or absence) of other studies that might contradict or qualify these findings but have been overlooked. It is worth noting that, unlike in cases typically discussed by Goldman (2001), the experts here are not contradicting each other. To be sure, their findings 'point' in different directions: while the epidemiological studies conclude that there is no proof of toxicity, the studies on rats provide grounds for suspecting it in humans. Nonetheless, these positions are entirely compatible. Moreover, the credibility of the experts themselves is not in question; no one disputes the methodological flaws in the epidemiological studies. However, these flaws do limit the relevance of the studies to the questions the judge must answer. The same applies to the studies on rats: what is disputed is their relevance to the issue of human toxicity. Thus, the task at hand is not, in principle, to choose between competing expert testimonies based on trustworthiness.

3.2.2 Should scientists speak to the relevance of the questions they are asked – and of the information they provide?

Suppose the judge's primary concern is to obtain impartial confirmation of the results observed in rats. They call on another toxicologist and ask the following question: *What do we know about the toxicity of the compound in rats?*[21]

Let us take the perspective of the expert. After confirming that experimental studies have observed serious cardiac side effects in rats given a certain dose of the compound over a certain period of time, should the expert go further and add that this information is of limited relevance to the human case – or specify the extent to which it may be relevant? Would that be overstepping their role?

Let us reverse the question: would an expert who limits themselves strictly to the terms of the question, and says nothing about the limited relevance of rat studies to human toxicity, be failing in their duty? Even if not privy to the full details of the case, the expert called to testify cannot ignore the fact that the trial concerns alleged human fatalities caused by the compound. In this sense, the expert is well aware that the central issue is human toxicity. And the judge knows that the expert knows this. For this reason, if the expert does not indicate that caution is required when extrapolating results from rats to humans, there is a real risk that their silence will be interpreted as a tacit endorsement of the inference to human toxicity.

Compare this to the case of a murder victim's roommate who is asked whether they heard the suspect's voice in the next room. Suppose they answer affirmatively, without mentioning that the voice came from their roommate's phone speaker. If later criticised for misleading the judge into concluding that the suspect was physically present in the next room, they might respond that they assumed the judge knew the voice came from a phone and that the question was about whether the victim and the suspect had spoken via speakerphone. Regardless of the sincerity of this response, we may argue that the judge is in a sufficiently informed position to have considered such a possibility; they are no less capable than the witness of realising that voices can be clearly heard through a phone speaker. In fact, being in a position of epistemic authority relative to the witness, the judge arguably should have phrased their question more precisely.

By contrast, when it comes to assessing the transposability of rat study results to humans, the judge is clearly less well equipped than the toxicology expert. It is even conceivable that the judge will not think to raise the issue, assuming instead that the findings in rats are probative of human toxicity. Therefore, if the judge fails to ask the expert what these results imply for our knowledge of the compound's effects in humans, it appears to fall to the expert to clarify the matter.

This example suggests that the expert is responsible for addressing at least certain aspects of the relevance of their statements for the judge. They must, in a sense, involve themselves in the inferences that the judge might draw from their testimony, insofar as the legitimacy of those inferences depends on the scientific standing of the claims made – that is, not only on their degree of confirmation but also on their inferential connections to other claims.[22]

Does this act of exceeding the question strictly posed, and entering the judge's domain of inference, violate the requirement of neutrality?[23] This requirement is grounded

21 We present here a largely idealised and simplified picture of scientists' testimony in a criminal trial, setting aside the distinction between reports submitted during the pre-trial phase and oral reports during the hearings. In addition, we assume – admittedly far from reality – a situation in which a single question is posed to the witness by the judge. These simplifications nonetheless serve to highlight issues that arise in real-world contexts.

22 It is worth emphasising that the task of clarifying the relevance of one's statements requires an effort from the expert that is unusual in the context of academic research. On the one hand, knowing how to distinguish between the various components of a scientific report, to assess their relative relevance and logical structure, is part of the scientist's craft. In reports intended for other scientists, such elements are generally left implicit, as they are taken for granted by those with similar academic training. On the other hand, in order to foster sound inferences – and to preempt flawed ones – the expert must anticipate the inferences the interlocutor is likely to draw from various possible responses, and thus adopt the judge's epistemic perspective.

23 We use the term neutrality *here in a non-technical sense, to refer simply to a duty of restraint on the part of the expert with respect to the decisions that the interlocutor will make on the basis of the information provided. Our analysis therefore does not intersect with philosophical debates about the possibility of 'value-free' scientific knowledge – that is, knowledge devoid of non-epistemic values (Longino 1996; Kitcher 2001, 2011; Douglas 2000; Rudner 1953; Wilholt 2009; Betz 2013; John 2015b; Parker 2014; Winsberg 2018).*

in a democratic ideal of informed decision-making: the expert supplies the most accurate factual information possible, and the decision-maker (individual or collective) then determines, on that basis, which course of action best advances their goals (Anderson 2011). On one interpretation of this ideal, the expert's neutrality requires that they limit themselves to factual responses and offer no guidance as to how the information might be used.[24]

In the cases of interest here, however, it is precisely the risk that the judge may draw faulty conclusions from the expert's testimony that makes it incumbent on the expert to issue a warning. 'Faulty conclusions' here means inferences that do not adhere to the rules of scientific reasoning, which demand that conclusions be grounded strictly in the current state of knowledge – a body of propositions that is not necessarily accessible to the judge. Such errors could lead to factual mistakes in the judge's reasoning. For this reason, it is to support the judge's 'epistemic well-being'[25] that the expert must speak to the relevance of their remarks. The aim is not to direct the judge toward a particular conclusion that the expert prefers, but to assist the judge in reaching a valid conclusion independently, without misinterpretation.

This task of epistemic facilitation requires the expert to accompany their responses with an assessment of their relevance. Doing so does not compromise their neutrality. A scientist who merely provides 'dry' answers to questions does not offer good expert testimony, in the sense that they fail to place the judge in the best epistemic position for decision-making.[26] The same holds, as we shall now see, for a scientist who provides information that, although accurate and technically responsive to the question, is not actually relevant.

3.2.3 Should the scientist be exhaustive?

Let us return to our example and now imagine that the judge seeks an impartial opinion from a toxicology specialist on the cardiac toxicity of the compound in question for humans. Without informing the expert of any prior testimony, the judge simply asks: *What is known about the cardiac toxicity of the compound in humans?* Should the specialist mention the existing data on rats?

If we adopt the academic criteria for what constitutes a good state of the art – as established by scientists for their peers (in research articles, review papers, grant proposals, or any other document internal to the academic sphere) – there would seem to be little doubt. A proper scientific response to a scientific question is an exhaustive one: it accounts for all relevant published material, meaning all sources that meet the standards of scientific credibility. Experimental data on a compound's toxicity in rats are part of what is scientifically relevant – insofar as they form part of what any scientist investigating the compound's human toxicity ought to know. While these data do not warrant conclusions about human toxicity, they can help orient and inform the investigation.[27] Indeed, animal testing is typically conducted precisely in the context of investigating the effects of substances on humans, as mandated by law.

However, from the judge's perspective, the relevant criterion for evaluating the expert's response is not whether it provides a comprehensive scientific overview. What matters for the judge is what they can draw from the expert's answer to the epistemic enquiry they are conducting in the course of the trial. In this context, the quality of an expert report must be assessed in terms internal to this legal-epistemic investigation. For the judge, the relevant answer to the question *What do we know about the cardiac toxicity of the compound in humans?* is above all one that can assist their reasoning on related matters – such as the liability of the pharmaceutical company. From that perspective, the rat studies are not relevant. Mentioning them could lead the judge to attribute to them a relevance they do not have, and to conclude – wrongly – that the compound is toxic in humans.

We may therefore conclude that the kind of exhaustiveness expected in academic science, where it serves as a criterion of credibility, is not always appropriate in a courtroom. Not only can the inclusion of information that is irrelevant to the

24 Except, of course, in cases where the question explicitly concerns how to achieve a given goal.

25 We borrow this expression from Stephen John (2018).

26 Our focus here differs from that of philosophers of science who, following the tradition of debates on inductive risk and value-free science (see note 23), view the irreducible uncertainty of science as a limit to the communication of neutral results. They ask, for instance, whether the standards for accepting scientific hypotheses should be sensitive to the context in which they are communicated (John 2015a; Steele 2012), and whether it is acceptable – or even desirable – for certain non-epistemic values to shape how scientists communicate findings that are fraught with uncertainty in public decision-making contexts, such as those concerning climate change. In that context, the issue arises as follows: given that scientific results often carry significant uncertainty, yet scientists must still provide answers to support decision-makers, they are effectively required to assume the risk inherent in any assertion. They must therefore determine – on the basis of non-epistemic values – the level of confirmation at which a hypothesis can be asserted. In the case we are discussing here, however, the issue is not necessarily one of uncertainty or values, but of the boundary between factual information and inference.

27 For example, certain results may prompt investigation into whether a compound follows a similar metabolic pathway in humans. Even when findings in rats are negative, research into potential human toxicity must still proceed – albeit in a more exploratory manner.

judge (even if scientifically accurate) confuse or complicate the message, it can actively mislead.[28] Once again, it is with respect to the communicative goal of promoting the epistemic well-being of the judge that the expert not only may, but arguably must, depart from a traditional scientific norm of credibility – namely, exhaustiveness.[29]

4. Conclusion

For a decision-maker seeking insight from a scientist, what matters is what the decision-maker can draw from the scientist's answers in relation to their own enquiry – not the standing of those answers within the broader body of scientific knowledge.[30] As a result, sound expertise for decision-making does not conform to the same criteria as a well-formed scientific report written for an academic audience. A report that aims to be a comprehensive review of the literature, without indicating the relative relevance of its contents, is likely not only to confuse and waste the time of a non-expert reader, but also to mislead – by suggesting false leads or encouraging unwarranted inferences. By contrast, good expertise consists in a report that both selects the scientific information relevant to the decision-maker (leaving out the rest) and presents it in a way that clarifies its relevance, thereby helping to guide the decision-maker along appropriate inferential paths.

This insight casts a new light on the demand that experts report as accurately as possible what science has to say on a given issue without interfering with the decision-making process. Applied too rigidly, this requirement can lead to communicative failure. While the expert must certainly not shape their response in order to promote the decision they themselves prefer, they must nonetheless help place the decision-maker in the best possible epistemic position to reach the right decision – in light of the decision-maker's own goals and values. To that end, the expert may sometimes need to indicate explicitly which inferences are or are not warranted based on the scientific claims presented. For similar reasons, they may also need to forgo exhaustiveness and omit certain information that could mislead the decision-maker simply by being mentioned. Indeed, as with any testimony, the act of presenting information also implicitly signals its relevance; offering irrelevant information in the name of completeness would jeopardise communication.

Unlike in the case of lay testimony, we have seen that the task of making the state of the art relevant seems to fall primarily to the expert, by virtue of their epistemic superiority with respect to the issues at hand. They alone are in a position to know which inferences are legitimately supported by what science can currently say. However, the task of clarifying the decision-maker's question cannot rest with scientists alone, and the idea that they bear sole responsibility for successful communication with non-scientists must be substantially qualified.

To offer sound expertise of the kind described here, the scientist must consider what motivates the question and what the decision-maker intends to do with the answer. Yet questions posed by non-scientific decision-makers are often 'ill-formed' by academic standards, and their motivations are not always clear-cut. A single question may serve multiple roles in a complex epistemic enquiry, and it may intersect with many other questions – some from other scientific disciplines, some from outside the scientific domain altogether.

This is especially evident in crisis situations, such as the COVID-19 pandemic, where urgency, high stakes, and widespread uncertainty are compounded by the need to integrate contributions from multiple disciplines – and to assess their relative relevance. In such a context, a scientist who lacks a detailed and comprehensive view of the problem, or of the decision-makers' implicit assumptions, values, and background knowledge, risks providing answers that are misaligned with what is actually needed.

To minimise that risk and foster productive communication between scientists and decision-makers within complex epistemic enquiries, the clarification of questions should be pursued through open dialogue – aimed at making explicit the assumptions on both sides, and at identifying the relative relevance of the various questions, as seen from both the scientists' and the decision-makers' standpoints.

28 This picture can be nuanced by suggesting that the appropriate response would be to mention the studies on rats, while clearly stating that their results cannot be extrapolated to humans. Among other things, this would protect the expert from the accusation of incompleteness, should the judge already be aware of such data. Clearly, the expected level of exhaustiveness also depends on the broader context of the request for expertise, and on the nature (and length) of the expected report. Nevertheless, presenting irrelevant information – alongside commentary on its limited relevance – may have the unintended effect, rather than enhancing credibility, of suggesting uncertainties beyond those already identified and scientifically measured, thereby undermining both the perceived credibility of science and the effectiveness of communication.

29 It thus seems that we may add exhaustivity and (a certain conception of) non-interference to the list proposed by Stephen John (2018) of virtues – such as transparency, openness, sincerity, and honesty – that are traditionally associated with good science communication, but which may in fact undermine it by failing to promote the epistemic well-being of the public.

30 The very notion of 'scientific knowledge' is therefore equivocal, depending on whether it refers to the current state of a question in the scientific literature – including uncertain hypotheses and weak data – or whether it is understood from outside the scientific community, where it is taken to embody a set of stable and robust claims.

REVUE
DE LA SOCIÉTÉ
DE PHILOSOPHIE
DES SCIENCES

What is good scientific expertise for decision-making? Relevance matters

BIBLIOGRAPHY

Anderson, Elizabeth. 2011. Democracy, Public Policy, and Lay Assessments of Scientific Testimony. *Episteme* 8 (2): 144-64. https://doi.org/10.3366/epi.2011.0013

Audi, Robert. 1997. The Place of Testimony in the Fabric of Knowledge and Justification. *American Philosophical Quarterly* 34 (4): 405-22.

Betz, Gregor. 2013. In defence of the value free ideal. *European Journal for Philosophy of Science* 3 (May). https://doi.org/10.1007/s13194-012-0062-x

Blais, Michel J. 1987. Epistemic Tit for Tat. *The Journal of Philosophy* 84 (7): 363. https://doi.org/10.2307/2026823

Bouchard, Frédéric. 2016. The Roles of Institutional Trust and Distrust in Grounding Rational Deference to Scientific Expertise. *Perspectives on Science* 24 (5): 582–608. https://doi.org/10.1162/POSC_a_00224

Bovens, Luc, and Stephan Hartmann. 2003. *Bayesian epistemology*. Oxford: Oxford; New York: Clarendon; Oxford University Press.

Burge, T. 1993. Content preservation. *The Philosophical Review* 102 (4): 457-88.

Coady, C. A. J. 1995. *Testimony: A Philosophical Study*. Oxford University Press, USA.

Collins, Harry, and Robert Evans. 2009. *Rethinking Expertise*. Paperback edition. Chicago London: University of Chicago Press.

Cook, R., I. W. Evett, G. Jackson, P. J. Jones, and J. A. Lambert. 1998. A Model for Case Assessment and Interpretation. *Science & Justice* 38 (3): 151-56. https://doi.org/10.1016/S1355-0306 (98)72099-4

Douglas, Heather. 2000. Inductive Risk and Values in Science. *Philosophy of Science* 67 (4): 559-79. https://doi.org/10.1086/392855

Dretske, F. I. 1981. *Knowledge and the Flow of Information*. Blackwell.

Fenton, Norman, Daniel Berger, David Lagnado, Martin Neil, and Anne Hsu. 2014. When 'Neutral' Evidence Still Has Probative Value (with Implications from the Barry George Case). *Science & Justice* 54 (4): 274-87. https://doi.org/10.1016/j.scijus.2013.07.002

Fricker, Elizabeth. 1995. Critical notice: Telling and trusting: Reductionism and anti-reductionism in the epistemology of testimony. *Mind* 104 (414): 393–411.

Goldman, Alvin I. 2001. Experts: Which Ones Should You Trust? *Philosophy and Phenomenological Research* 63 (1): 85–110. https://doi.org/10.2307/3071090

Grice, Herbert Paul. 1995. *Studies in the Way of Words*. 4. print. Cambridge, Mass: Harvard Univ. Pr.

Hahn, Ulrike, Adam J. L. Harris, and Adam Corner. 2009. Argument Content and Argument Source: An Exploration. *Informal Logic* 29 (4): 337. https://doi.org/10.22329/il.v29i4.2903

Hardwig, John. 1985. Epistemic Dependence. *The Journal of Philosophy* 82 (7): 335-49.

Hardwig, John. 1991. The Role of Trust in Knowledge. *The Journal of Philosophy* 88 (12): 693–708. https://doi.org/10.2307/2027007

Hempel, Carl G. 1960. Science and Human Values. In *Social Control in a Free Society*, edited by Robert E. Spiller, 39–64. University of Pennsylvania Press. https://doi.org/10.9783/9781512807424-003

Jeffrey, Richard C. 1996. *The Logic of Decision*. 2. ed. Chicago: Univ. of Chicago Press.

John, Stephen. 2015a. Inductive risk and the contexts of communication. *Synthese* 192 (1): 79–96. https://doi.org/10.1007/s11229-014-0554-7

John, Stephen. 2015b. The example of the IPCC does not vindicate the Value Free Ideal: a reply to Gregor Betz. *European Journal for Philosophy of Science* 5 (1): 1–13. https://doi.org/10.1007/s13194-014-0095-4

John, Stephen. 2018. Epistemic trust and the ethics of science communication: against transparency, openness, sincerity and honesty. *Social Epistemology* 32 (2): 75–87. https://doi.org/10.1080/02691728.2017.1410864

Kitcher, Philip. 2001. *Science, Truth, and Democracy*. Oxford University Press. https://doi.org/10.1093/0195145836.001.0001

Kitcher, Philip. 2011. *Science in a democratic society*. Amherst, N.Y: Prometheus Books.

Lagnado, David A., and Tobias Gerstenberg. 2017. 'Causation in Legal and Moral Reasoning.' In *The Oxford handbook of causal reasoning*, edited by Michael Waldmann, 565–601. Oxford library of psychology. New York, N.Y.: Oxford University Press.

Lagnado, David A., and Nigel Harvey. 2008. The Impact of Discredited Evidence. *Psychonomic Bulletin & Review* 15 (6): 1166-73. https://doi.org/10.3758/PBR.15.6.1166

Longino, Helen E. 1996. Cognitive and Non-Cognitive Values in Science: Rethinking the Dichotomy. In *Feminism, Science, and the Philosophy of Science*, edited by Lynn Hankinson Nelson and Jack Nelson, 39–58. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-009-1742-2_3

Parker, Wendy. 2014. Values and Uncertainties in Climate Prediction, Revisited. *Studies in History and Philosophy of Science Part A* 46 (June): 24–30. https://doi.org/10.1016/j.shpsa.2013.11.003

Roberts, Paul. 2022. *Roberts & Zuckerman's criminal evidence*. Third edition. Oxford, United Kingdom: Oxford University Press.

Rudner, Richard. 1953. The Scientist Qua Scientist Makes Value Judgments. *Philosophy of Science* 20 (1): 1–6.

Sackett, D. L, W. M C Rosenberg, J A M. Gray, R B. Haynes, and W S. Richardson. 1996. Evidence Based Medicine: What It Is and What It Isn't. *BMJ* 312 (7023): 71–72. https://doi.org/10.1136/bmj.312.7023.71

REVUE
DE LA SOCIÉTÉ
DE PHILOSOPHIE
DES SCIENCES

Saul, Jennifer Mather. 2012. *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199603688.001.0001

Schneps, Leila. 2018. Statistics, probabilities and justice. *Cahiers philosophiques* N° 155 (4): 21. https://doi.org/10.3917/caph1.155.0021

Schum, David A. 1994. *The Evidential Foundations of Probabilistic Reasoning*. Evanston, Ill.: Northwestern University Press.

Solan, Lawrence, and Peter Meijes Tiersma. 2005. *Speaking of crime: the language of criminal justice*. The Chicago series in law and society. Chicago: University of Chicago Press.

Sperber, Dan, and Deirdre Wilson. 2001. *Relevance: communication and cognition*. 2nd ed. Oxford; Cambridge, MA: Blackwell Publishers.

Steele, Katie. 2012. The Scientist qua Policy Advisor Makes Value Judgments. *Philosophy of Science* 79 (5): 893–904. https://doi.org/10.1086/667842

Stegenga, Jacob. 2018. *Medical nihilism*. First edition. Oxford, United Kingdom: Oxford University Press.

Thayer, James Bradley. 1969. *A preliminary treatise on evidence at the common law*. South Hackensack, N.J: Rothman Reprints.

Wilholt, Torsten. 2009. Bias and values in scientific research. *Studies in History and Philosophy of Science Part A* 40 (1): 92–101. https://doi.org/10.1016/j.shpsa.2008.12.005

Winsberg, Eric. 2018. Communicating Uncertainty to Policymakers: The Ineliminable Role of Values. In *Climate Modelling*, edited by Elisabeth A. Lloyd and Eric Winsberg, 381–412. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-65058-6_13

HISTORIQUE

Article soumis le 26 mai 2025.
Article accepté le 1 juin 2025.

SITE WEB DE LA REVUE

https://ojs.uclouvain.be/index.php/latosensu

DOI

https://doi.org/10.20416/LSRSPS.V12I1.3

CONTACT ET COORDONÉES

Anouk Barberousse
Sorbonne-Université
SND, Bureau G071
17 Rue de Sorbonne
75005 PARIS

Marion Vorms
Université Paris 1
IHPST
13 rue du Four
75006 PARIS

SOCIÉTÉ DE PHILOSOPHIE DES SCIENCES (SPS)

École normale supérieure
45, rue d'Ulm
75005 Paris

SOCIÉTÉ DE PHILOSOPHIE DES SCIENCES (SPS)

École normale supérieure
45, rue d'Ulm
75005 Paris
www.sps-philoscience.org