★

Towards Science-Friendly Epistemology

★

Jon Williamson
Philosophy Department, University of Manchester

in Y. Shan (ed.), *Philosophy of Science and Epistemology*, Springer

Abstract

The gap between philosophy of science and epistemology is both wide and widely bemoaned. But what should science-friendly epistemology (SFE) look like? In this paper I put forward and discuss five desiderata. (i) SFE should be grounded in real scientific practice: it should be motivated by and able to cope with the complexities of real science. (ii) SFE should avoid 'metaphysical bias', i.e., metaphysical presuppositions that constrain the search space of epistemological theories, potentially leading to implausible claims about science or sub-optimal scientific methods. Yet it should not be merely descriptive: (iii) SFE should offer the potential to improve scientific practice. On the other hand, it shouldn't be too revisionary with respect to the conceptual core of science: (iv) SFE should generally appeal to concepts with which scientists are familiar and use them in the ways in which they are used in science. Finally, (v) SFE should focus more on cases that are typical or important in science and less on far-fetched putative counterexamples.

I then consider Evidential Pluralism as a potential exemplar of science-friendly epistemology. Evidential Pluralism is a specific theory of the epistemology of causality: according to Evidential Pluralism, establishing causation requires establishing both the existence of a correlation and the existence of a mechanism, and assessing causation requires assessing mechanistic studies alongside association studies. After introducing this epistemological theory, I argue that it meets all five desiderata for SFE.

§1
Bridging the gap

One would be forgiven for thinking that epistemology, which seeks to understand how we can obtain knowledge, and philosophy of science, which seeks to understand our best available methods for obtaining knowledge, should be close bedfellows. It is thus surprising how far the two fields have diverged over the years. As Otávio Bueno notes, for instance,

> It is a sad fact of contemporary epistemology and philosophy of science that there is very little substantial interaction between the two fields. Most epistemological theories are developed largely independently of

> any significant reflection about science, and several philosophical interpretations of science are articulated largely independently of work done in epistemology. (Bueno, 2016, p. 233.)

In response to this divergence, there have been some isolated attempts to reconcile the two fields. For example, Bird (2022) develops an anti-empiricist philosophy of science by appealing to developments in knowledge-first epistemology; Ludwig et al. (2023) seek to reconceive the two fields as much more diverse, interconnected and globally engaged than they might seem; and Berghofer (2025) connects internalist epistemology with quantum Bayesianism. The areas of formal epistemology (e.g., Bradley, 2015) and naturalised epistemology (e.g., Schmitt, 2004) also provide points of interaction between epistemology and philosophy of science.

In this paper, I attempt to bridge the gap in a more general way by asking how epistemology might become more relevant to science and scientists. I put forward a vision for what I call 'science-friendly epistemology,' which can be thought of as being situated in the intersection of epistemology and philosophy of science, and I suggest that both epistemology and science could benefit if a greater proportion of research in epistemology were science-friendly.

What kind of epistemological theory would qualify as 'science-friendly'? In §2 I develop five desiderata that such a theory should satisfy. These desiderata are quite demanding, and one might reasonably wonder whether it is possible to satisfy all five. I shall argue that Evidential Pluralism, an emerging theory of causal enquiry (see §3), satisfies all five desiderata and can therefore be viewed as an exemplar of science-friendly epistemology (§4).

§2
Desiderata for Science-Friendly Epistemology

The aim of this section is to provide an account of what constitutes science-friendly epistemology ('SFE' for short). In order to pin down what I mean by SFE, I shall take SFE to have the following overarching goal: *to develop epistemological theories that provide useful accounts of scientific enquiry.* I shall argue that this goal motivates five desiderata. A science-friendly epistemological theory should: (i) be grounded in current scientific practice; (ii) avoid metaphysical presuppositions that may lead to sub-optimal epistemological hypotheses; (iii) offer the potential to improve scientific practice; (iv) generally appeal to concepts with which scientists are familiar and use them in the ways in which they are used in science; (v) focus more on cases that are typical or important in science and less on far-fetched putative counterexamples.

Let us first clarify the overarching goal itself, i.e., the characterisation of SFE as seeking epistemological theories that provide useful accounts of scientific enquiry. By 'useful,' I mean useful to scientists themselves—SFE should offer theories that can help scientists with their enquiry. This might mean providing explicit methods for scientific enquiry, or it could mean providing helpful ways of thinking about scientific enquiry. By 'scientific enquiry,' I mean establishing and evaluating scientific claims. Epistemology usually focusses on epistemic concepts, such as knowledge and belief, that are subjective in the sense that they are predicated of individual subjects or groups. While the knowledge and beliefs of scientists and research groups may be of some interest to epistemologists and philosophers of science, scientists themselves tend to focus on more impersonal epistemic tasks, such as establishing and evaluating claims. Epistemology has recently begun to recognise the need to theorise about enquiry:

> epistemology as we know it, as it is currently done, doesn't seem to be leaving as much room as we might like for the sorts of norms that are central to the practice of inquiry. (Friedman, 2020, p. 502.)

Philosophy of science, on the other hand, has for some time offered accounts of enquiry that scientists find useful. These include very general accounts of scientific enquiry, such as Popper's falsificationism (Popper, 1934) and Bayesian confirmation theory (Howson and Urbach, 1989), as well as accounts that are specific to particular kinds of enquiry—accounts of causal enquiry, for example.

I should emphasise that my suggestion is that more of epistemology should be science-friendly, not that all of epistemology should be science-friendly. The questions of orthodox epistemology remain, and remain of interest. Indeed, SFE could be as helpful for epistemology as it could be for science. By bridging the gap between epistemology and philosophy of science, SFE could help to invigorate epistemology by facilitating the flow of new ideas from philosophy of science to epistemology. Likewise, SFE could help to facilitate the flow of ideas from epistemology to the philosophy of science. This sort of bidirectional flow of ideas between two autonomous fields—what Gillies and Zheng (2001) call a 'dynamic interaction'—can be of substantial mutual benefit to the two fields (see §5 for an example).

⋆

Having said something about what I mean by 'science-friendly epistemology', I shall now put forward five desiderata for a science-friendly epistemological theory.

D1. Real Science. A science-friendly epistemological theory should be grounded in current scientific practice.

Recall the overarching goal of SFE: to develop epistemological theories that provide useful accounts of scientific enquiry. It is hard to see how an epistemological theory could offer a useful account of scientific enquiry if it were incompatible with, or even neutral with respect to, real scientific enquiry. The first desideratum for a science-friendly epistemological theory is thus that it should be motivated and confirmed by current science. One should be able to appeal to scientific practice to see both that the theory is plausible and why it is plausible. Moreover, a science-friendly epistemological theory should be able to cope with the complexities of real science. Epistemology is already replete with theories that are grounded in toy examples but offer little insight as to their consequences for scientific enquiry in all its complexity. To be useful to scientific enquiry, SFE should have clear consequences for that enquiry in all its complexity.

Some epistemological theories are grounded in metaphysical claims, rather than in current scientific practice. Metaphysical claims have some heuristic value in that they can suggest new epistemological hypotheses. For example, a counterfactual view of causality can be used to motivate interventionist approaches to causal enquiry or formalisms for causal enquiry such as the potential outcomes approach (see, e.g., Rubin, 2005). But by constraining the search space of epistemological hypotheses, metaphysical claims also constitute *biases*: they can lead to false hypotheses and sub-optimal methods (see §4.5 on this point). Andersen et al. (2019) call such biases 'philosophical biases':

> Basic philosophical assumptions count as biases because they skew the development of hypotheses, the design of experiments, the evaluation of evidence, and the interpretation of results in specific directions. (Andersen et al., 2019, p. 1.)

Metaphysical bias, then, constitutes one kind of philosophical bias.

Now, scientists try to identify and eliminate biases where possible. While heuristics and biases help us to 'think fast' in everyday life, successful science requires thinking slowly, trading off speed for reliability of inferences. This leads to our second desideratum:

D2. Metaphysical Bias. A science-friendly epistemological theory should avoid metaphysical bias, where possible.

It is worth noting that, in advocating D2, I depart from Andersen et al., who suggest that philosophical biases cannot be avoided:

> To give an example related to causality: when choosing a scientific method to establish a causal relationship between some medical condition and a virus, one must first have an idea of what causality is. This is a part of science that cannot be discovered empirically, but remains tacitly assumed in scientific methodology and practice. (Andersen et al., 2019, p. 1.)

In §4 I shall argue against this claim of Andersen et al.: I shall show that one does *not* first need to have an idea of what causality is when providing an account of causal enquiry. Since it is possible to avoid this metaphysical bias, one ought to avoid it, according to D2. Of course, one cannot avoid *all* metaphysical claims: it is hard, for example, to see how one can have a coherent account of the role of experiment in science without assuming that the external world exists. Where possible, however, it is important for SFE to identify and mitigate metaphysical bias in order to reduce the chance of settling on sub-optimal epistemological hypotheses about science. D2, then, goes well beyond the position of Andersen et al., who focus on articulating and exploring metaphysical biases rather than avoiding them.[1]

So far, then, our desiderata have concerned the direction of motivation for SFE, putting forward what might be called a 'science-first' approach: a science-friendly epistemological theory should be motivated by current science, rather than constrained by metaphysics. Our next desideratum is that SFE should extend beyond any motivating examples:

D3. Normativity. A science-friendly epistemological theory should offer the potential to improve scientific practice.

In the light of our guiding principle, the motivation behind this desideratum is straightforward: if a theory were merely descriptive, it is hard to see how it could

[1]Andersen et al.'s approach has elements in common with that of Collingwood (1940), who sees an important role in uncovering the metaphysical presuppositions of enquiry.

be of much use to scientists in informing aspects of scientific enquiry that extend beyond the cases within the scope of the descriptive theory.

A fourth desideratum also focusses on the usability of a theory:

D4. Scientific Concepts. A science-friendly epistemological theory should appeal to concepts with which scientists are familiar and use them in the ways in which they are used in science, where possible.

Theories that are very revisionary with respect to the conceptual core of science are less science-friendly because they force scientists to abandon entrenched ways of working and to use new concepts with which they lack expertise. Satisfying this desideratum may well be somewhat frustrating for philosophers, who focus almost exclusively on claims at the core of the Quinean web of beliefs and who are often esteemed in proportion to the extent to which they creatively challenge these core beliefs. But epistemological theories that fail to satisfy this desideratum can be equally frustrating for scientists. Of course, sometimes the conceptual core of science must change in order for scientific theories to remain empirically adequate in the light of new evidence. But this process is typically driven by empirical anomalies or theoretical inconsistencies, not by philosophical considerations. The move to wave-particle duality, for example, was arguably required to accommodate experimental evidence, not because wave-particle duality is somehow philosophically appealing.

Finally,

D5. Typicality. A science-friendly epistemological theory should focus more on cases that are typical or important in science and less on far-fetched putative counterexamples.

Fake barns may be of great interest to epistemologists (see, e.g., Colaço et al., 2014), but they are of much less immediate relevance to day-to-day scientific enquiry. In order to offer useful accounts of scientific enquiry, SFE needs to say something useful about the typical or important cases that form the bread and butter of scientific endeavour.

⋆

In sum, if the aim of SFE is to offer useful accounts of scientific enquiry then the following desiderata become plausible:

D1. Real Science. A science-friendly epistemological theory should be grounded in current scientific practice.

D2. Metaphysical Bias. A science-friendly epistemological theory should avoid metaphysical bias, where possible.

D3. Normativity. A science-friendly epistemological theory should offer the potential to improve scientific practice.

D4. Scientific Concepts. A science-friendly epistemological theory should appeal to concepts with which scientists are familiar and use them in the ways in which they are used in science, where possible.

D5. Typicality. A science-friendly epistemological theory should focus more on cases that are typical or important in science and less on far-fetched putative counterexamples.

But perhaps these desiderata are asking too much of an epistemological theory. It would be a problem for the project of SFE if no epistemological theory satisfied the desiderata. In order to show that SFE is non-empty, I introduce Evidential Pluralism in the next section and then argue in §4 that it can be viewed as an exemplar of SFE.

§3
Evidential Pluralism

Evidential Pluralism offers a philosophical account of causal enquiry, i.e., of how to establish and evaluate causal claims. It is an ongoing, collaborative research programme: the key ideas were put forward by Russo and Williamson (2007) and Illari (2011) and the account was subsequently developed in the contexts of the health sciences (see, e.g., Parkkinen et al., 2018; Gillies, 2019), the social sciences (Shan and Williamson, 2023), and law (Trofimov and Williamson, 2025). This line of research can be thought of as an epistemological strand of the 'new mechanism' approach to philosophy of science, which emphasises the importance of mechanisms to scientific enquiry (Craver et al., 2024). In this section, I shall very briefly summarise some of the key tenets of Evidential Pluralism; see Shan and Williamson (2023, Chapter 1) and Williamson (2019a) for fuller treatments of the theory and its motivation.

According to Evidential Pluralism, in order to establish that A is a cause of B one needs to establish (i) that A is correlated with B, conditional on potential confounders and (ii) there is some mechanism complex linking A to B that can account for the extent of the correlation. Consequently, in order to assess whether A is a cause of B, one needs to scrutinise both association studies (which repeatedly measure A and B, often together with potential confounders, in order to assess the extent of any association) and mechanistic studies (which seek evidence of features of hypothesised mechanisms linking A and B, such as mediating variables, entities, activities and organisational structure).

The key evidential relationships posited by Evidential Pluralism are depicted in Fig. 1. According to Evidential Pluralism, establishing causation requires establishing the *existence* of a correlation and the *existence* of a mechanism, not the extent of the correlation nor the details of the mechanism. One can of course establish the existence of correlation by estimating the extent of the correlation (via channel α_1 in Fig. 1) and one can establish the existence of a mechanism by identifying key features of that mechanism (μ_1 and μ_2), but association studies can provide indirect evidence of the existence of a mechanism (α_2) and mechanistic studies can provide indirect evidence of the existence of a correlation (μ_3). Among association studies, randomised controlled trials (RCTs) are particularly valued because they provide indirect evidence of the existence of a mechanism (α_2), as well as direct evidence of the existence of a correlation (α_1): randomisation lowers the probability that any observed correlation is attributable to unforeseen confounding and thereby raises the probability that it is attributable to the existence of some mechanism by which A causes B.

These evidential relationships can be further extended by considering 'source studies', i.e., studies carried out on a different source population. In the right

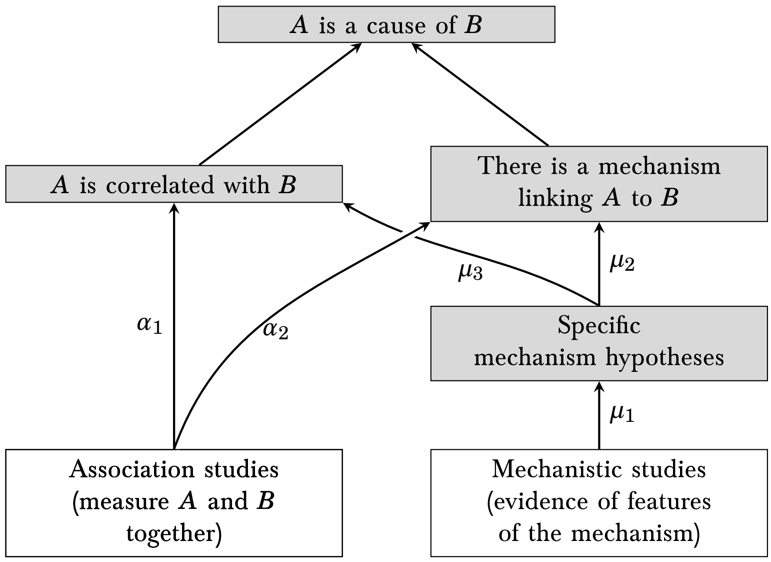

Figure 1: Evidential relationships when establishing and evaluating causal claims.

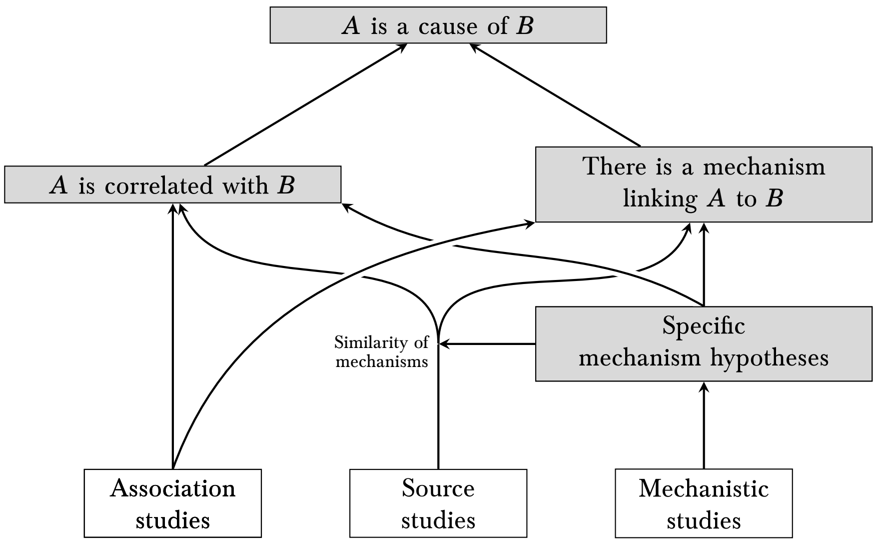

Figure 2: Evidential relationships including extrapolation from a source population.

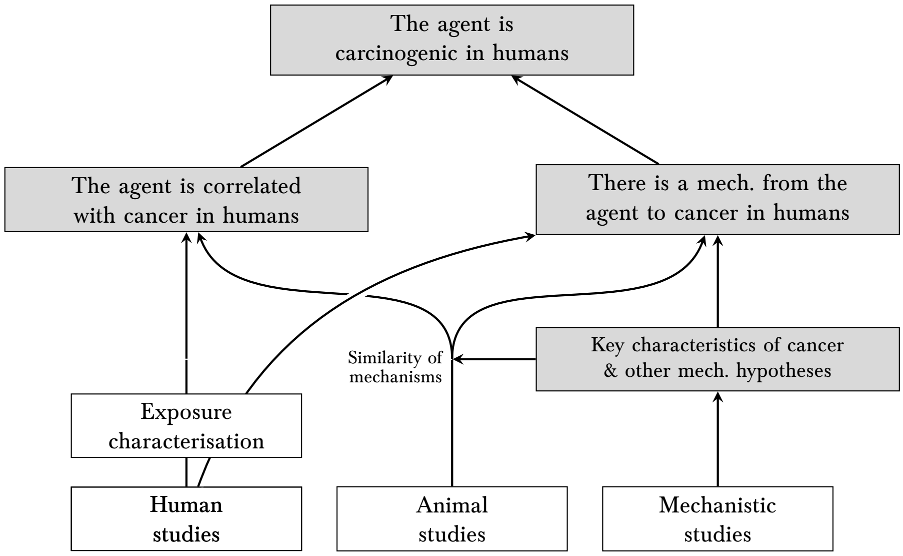

Figure 3: Evidential relationships in an IARC carcinogenicity evaluation.

circumstances, source studies can confirm or undermine the causal claim of interest in the target population of interest (Fig. 2). Here, Evidential Pluralism incorporates a version of mechanism-based extrapolation (Steel, 2008): evidence that A causes B in a source population is relevant to the claim that A causes B in the target population when the mechanism complexes underlying the causal relationship in source and target populations are sufficiently similar. Evidence of similarity can be provided by demonstrating that key features of the mechanism of action are the same or very similar in the source and target populations, that any enabling or enhancing mechanisms in the source population are also present in the target population, and that no new interacting mechanisms in the target population could fully counteract the influence of the mechanism of action there.

§4
Evidential Pluralism as an exemplar of SFE

In this section, I shall argue that Evidential Pluralism satisfies the desiderata of §2 to such an extent as to count as 'science-friendly'.[2]

§4.1. Grounding in real scientific practice

When motivating the leading idea behind Evidential Pluralism, Russo and Williamson (2007) explicitly appealed to three aspects of scientific practice: methodological

[2]Note that Evidential Pluralism is an account of specifically *causal* enquiry. Causal enquiry is ubiquitous in the biomedical and social sciences and in at least some parts of the physical sciences, but I do not claim that all scientific enquiry is causal enquiry. Many argue, for example, that non-causal explanation is an important kind of scientific enquiry (e.g., Lange, 2017). Other epistemological theories would be needed to address these other kinds of enquiry. This yields a piecemeal approach to scientific enquiry. There is nothing inherently problematic about a piecemeal approach, however. We carve up different kinds of enquiry precisely because they require differences in approach.

Table 1: Examples of IARC Monographs classifications based on mechanistic evidence; based on Benbrahim-Tallaa (2019). 2A = probably carcinogenic to humans; 3 = not classifiable as to its carcinogenicity to humans.

Agent	Mechanistic evidence	Eval.
d-Limonene	The mechanism by which *d*-limonene increases the incidence of renal tubular tumours in male rats is not present to humans	3
Simian virus (SV)40	No persuasive evidence that the mechanism of transformation in rodents is operative in humans	3
1-Nitropyrene	Is genotoxic, induces oxidative stress	2A
1,3-Propane sultone	Is genotoxic	2A
Tetrabromobisphenol A	Modulates receptor-mediated effects, induces oxidative stress, is immunosuppressive	2A
3,3',4,4'-Tetrachloroazobenzene	Belongs to a class of agents for which one or more members have been classified as carcinogenic or probably carcinogenic to humans	2A

practice, case studies of causal enquiry, and the goals of causal enquiry. Let us consider these three aspects in turn.

Firstly, with regard to methodological practice, Russo and Williamson (2007) argued that Austin Bradford Hill's criteria for establishing causality can be underpinned by Evidential Pluralism, as can the methodology of the carcinogenicity evaluations of the International Agency for Research on Cancer (IARC). Let us consider IARC methodology. When evaluating a potential cause of cancer in humans (e.g., an exposure to a particular chemical in the workplace), IARC convenes a group of experts and divides it into subgroups. One subgroup evaluates epidemiological studies on humans—i.e., association studies—, while another evaluates mechanistic studies, and a third subgroup assesses studies carried out on experimental animals—i.e., source studies. Importantly, IARC's 'Preamble', which encapsulates its evaluation methodology, has recently been revised to explicitly treat mechanistic studies on a par with epidemiological studies (IARC, 2019; Samet et al., 2020):

> The revision of the Preamble took into account advances in the assessment of mechanistic data, including, in particular, the identification of key characteristics of carcinogens, which provide a framework for the organization of mechanistic data and the assessment of strengths as well as gaps in evidence. The current Preamble reflects these advances and describes a process to reach a carcinogenicity classification by integrating, along parallel and harmonized lines, the three streams of evidence: experimental animal bioassays, mechanistic investigations, and epidemiological studies. (Berrington de González et al., 2024, p. 208.)

Table 1 provides some examples of IARC evaluations in which mechanistic evidence played a prominent role. In all these examples, human studies provided inadequate evidence, but carcinogenicity in animal models was established. In the cases of d-Limonene and Simian virus, mechanistic evidence undermined the claim that the agents are carcinogenic in humans, because it undermined the claim that the animal carcinogenicity mechanisms are present in humans. In the other examples, mechanistic evidence led to greater confidence in carcinogenicity than would otherwise be warranted, because key features of carcinogenicity mechanisms were found to be present in humans. Thus, IARC's evaluation procedure can be viewed as conforming very closely to Evidential Pluralism, as depicted in Fig. 3 (Williamson, 2019b).[3] The success of the IARC carcinogenicity evaluation programme suggests that Evidential Pluralism is essentially correct in its claims about the evidential relationships in causal enquiry. It also shows that Evidential Pluralism can be successfully implemented in practice (Williamson, 2021).

Let us turn to the second aspect of scientific practice that Russo and Williamson (2007) invoke to motivate Evidential Pluralism, namely particular case studies of causal enquiry, including the case of Snow's discovery of the transmission of cholera, Semmelweis' claims about the prevention of puerperal fever, and the more recent case of the discovery that ***Helicobacter pylori*** is a cause of ulcers. Since 2007, many more instances of causal enquiry have been analysed and found to conform to Evidential Pluralism. For example, Clarke (2011) considered viral causes of cancer and also McArdle disease; Gillies (2019) considered atherosclerosis and coronary heart disease, smoking and lung cancer, and also streptomycin and tuberculosis; Maziarz and Stencel (2022) analysed hydroxychloroquine and COVID-19; Park et al. (2023) considered efavirenz and HIV; and Shan and Williamson (2023) considered a range of cases across the social sciences.

The third line of motivation for Evidential Pluralism is an appeal to the goals of scientific practice (Russo and Williamson, 2007, p. 159). Across the sciences, causal claims are used for prediction, explanation and control. The use of causal claims for prediction (predicting an effect from the presence of a cause, or diagnosing the cause from the effect) and control (manipulating a cause to change an effect) requires that a cause makes a difference to its effects, i.e., that cause and effect are correlated. On the other hand, the use of causal claims for explanation (explaining the occurrence of an effect by citing its causes) requires that there is some mechanism complex by which its causes are responsible for the effect, that can flesh out the explanation. Thus the various uses of causal claims in science also help to motivate the need to establish both correlation and mechanism when establishing causation.

Evidential Pluralism is thus grounded in three aspects of scientific practice—methodological practice, case studies, and the goals of scientific enquiry—, in line with desideratum D1:

[3] Iranzo and Pérez-González (2024, p. 9) suggest that the fit between IARC's evaluation procedure and Evidential Pluralism is imperfect, on the grounds that

> IARC evidential policy undeniably gives some priority to statistical evidence. For instance, concerning a crucial category as Group 1—i.e., "carcinogenic"—, evidence of mechanisms is not necessary to get a conclusive result, while statistical evidence in humans is indispensable.

It is true that IARC's evaluation procedure does not *require* high-quality mechanistic studies to establish carcinogenicity. But this does accord with Evidential Pluralism: association studies can establish causality via channels α_1 and α_2 of Fig. 1. So, this aspect of IARC's procedure should not be construed as out of line with Evidential Pluralism.

D1. Real Science. A science-friendly epistemological theory should be grounded in current scientific practice.

§4.2. Avoiding metaphysical bias

Recall Andersen et al. (2019, p. 1): 'when choosing a scientific method to establish a causal relationship between some medical condition and a virus, one must first have an idea of what causality is. This is a part of science that cannot be discovered empirically, but remains tacitly assumed in scientific methodology and practice.' Their suggestion is that an idea of what causality is amounts to a philosophical bias that cannot be avoided (but that nevertheless needs to be made explicit and open to criticism).

Evidential Pluralism can be invoked to refute their suggestion that it is impossible to avoid metaphysical bias when settling on a method for causal enquiry. This is because Evidential Pluralism does not presuppose any particular idea of what causality is. Instead, it is motivated by cases of real scientific practice, as we have just seen. The direction of motivation thus proceeds from science to epistemology: it is the need to accommodate real scientific practice, rather than the desire to accommodate some particular view of causality, that leads to Evidential Pluralism.[4]

In fact, Evidential Pluralism is compatible with a variety of markedly different metaphysical accounts of causality. Originally, Russo and Williamson (2007) used Evidential Pluralism to argue for an anti-realist account of causality, namely the ***epistemic theory of causality***, which analyses causal relationships in terms of rational belief (in a way that is analogous to the way in which Bayesianism analyses probabilities in terms of rational belief); see also Shan and Williamson (2023, §5). Weber (2009, 2025) subsequently argued that Evidential Pluralism is also compatible with Ronald Giere's metaphysical theory of causality, and Giere's theory reduces causal relationships to counterfactual correlations. In addition, Evidential Pluralism is compatible with realist accounts of causality that take causation to be a primitive, irreducible relation in the world (e.g., Tooley, 1990; Chakravartty, 2005). That Evidential Pluralism is apparently compatible with realism, anti-realism and reductivism about causality indicates that it does not presuppose any such account.

This is not to say that Evidential Pluralism has no metaphysical consequences whatsoever. If Evidential Pluralism is correct, this clearly rules out those metaphysical accounts of causality that fail to validate Evidential Pluralism: a difference-making account of causation that requires only evidence of correlation, for example, or a mechanistic account that only sees a need for mechanistic evidence (see Fig. 4), or an eliminativist account of causality that sees a need for neither. Moreover, Evidential Pluralism appeals to notions of correlation and of mechanism, and a place is needed for these in any ontology that supports Evidential Pluralism. But Evidential Pluralism does not presuppose any particular metaphysical account of

[4]The cases of real scientific practice that are used to motivate Evidential Pluralism may themselves be prone to philosophical bias, because the scientists producing those studies may have preconceptions about the metaphysics of causality that influence their research. This does not imply that Evidential Pluralism itself is adversely affected by metaphysical bias, however. This is because different motivating studies will be influenced by different (and contradictory) biases, and these biases can get washed out when generalising over cases to yield a science-friendly causal epistemology such as Evidential Pluralism. Consider an analogy. A meta-analysis aims by generalising over multiple studies to lessen the influence of the inevitable biases that affect the individual studies included in the meta-analysis. It can do this because different studies are influenced by different biases that often pull in different directions.

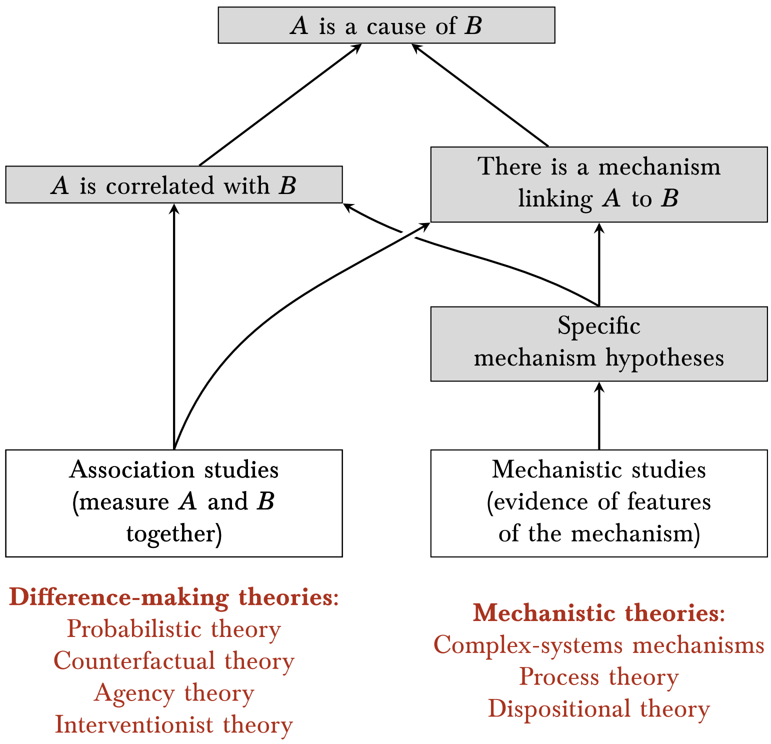

Figure 4: Metaphysical theories of causality may bias causal enquiry towards methods allied to one side of this diagram, erroneously excluding methods allied to the other side.

causation, correlation or mechanism, so it avoids metaphysical bias, in line with D2:

D2. Metaphysical Bias. A science-friendly epistemological theory should avoid metaphysical bias, where possible.

§4.3. Providing normative guidance

As Fig. 4 suggests, methods for causal enquiry that focus on association studies to the exclusion of mechanistic studies, or vice versa, can be sub-optimal. One way to reach a flawed methodology is by being guided by metaphysics, but another way is by taking a dogmatic attitude towards a particular set of research methods or evaluation methods—what I shall call 'methodological dogma'.

According to proponents of Evidential Pluralism, orthodox 'evidence-based' evaluation methods have become methodological dogma. Through their exclusive focus on association studies—particularly randomised controlled studies (RCTs) of the association between the putative cause and effect—orthodox evidence-based medicine and evidence-based policy tend to overlook mechanistic studies:

> Evidence-based medicine de-emphasizes intuition, unsystematic clinical experience, and pathophysiologic rationale as sufficient grounds for

> clinical decision making and stresses the examination of evidence from clinical research. (Guyatt et al., 1992, p. 2420.)

If Evidential Pluralism is correct, however, then mechanistic studies should be scrutinised alongside association studies such as RCTs. Thus, the dogmatic adoption of orthodox evidence-based evaluation has arguably led to the widespread use of sub-optimal evaluation methods.

One can diagnose the problem as follows. Orthodox evidence-based evaluation methods tend to conflate two questions: (i) In the absence of other relevant evidence, and if cost were no object, what sort of study should we carry out to evaluate whether an intervention is effective? (ii) What sorts of studies should we scrutinise in order to evaluate whether an intervention is effective? A common answer to (i) is that, if it is possible to properly randomise and blind the intervention to individuals in the target population of interest, then a large, well-conducted RCT is likely to be most informative (via channels α_1 and α_2 of Fig. 1). It is clearly fallacious to infer, however, that one should focus solely, or even mainly, on RCTs in order to answer (ii). The RCTs that are available may be poorly conducted, or insufficiently large, or improperly blinded, or improperly randomised, or randomised to groups rather than individuals (with the potential for confounding by group characteristics), or sampled from a population other than the one upon which we intend to intervene (perhaps by excluding certain kinds of individuals from the trials). Such problems are ubiquitous in the areas of medicine, public health and public policy, for example, and any one of these limitations is reason enough to broaden the evidence base. Even in the absence of these problems, there can be good reasons to broaden the evidence base. Even if the RCTs are all of unusually high quality, their results may point in different directions and thus fail support any overall conclusion.[5] Or it may be that credible mechanism hypotheses undermine the results of the available RCTs by pointing to possible sources of bias or confounding that have not been adequately controlled for. (Recall that RCTs do not eliminate the problem of confounding—they merely reduce its long-run.) So there is good reason to scrutinise observational studies and mechanistic studies alongside RCTs—particularly when the RCTs have limitations or are otherwise inconclusive or are undermined by mechanism hypotheses.

Evidential Pluralism thus motivates an alternative evaluation methodology—one that systematically scrutinises mechanistic studies in addition to the association studies that are the focus of orthodox evidence-based methods (see Table 2). This leads to a new approach to evidence-based medicine called 'EBM+' (Parkkinen et al., 2018), an analogous 'EBP+' approach to evidence-based policy (Shan and Williamson, 2023), and a corresponding 'EBL+' approach to evidence-based law and regulation (Trofimov and Williamson, 2025).

This new methodology has the potential to improve evaluation practice. Aronson et al. (2021), for example, argue that we should have been using this new methodology during the Covid-19 pandemic. Greenhalgh et al. (2022) concur:

> Our central argument is that for some aspects of the pandemic, especially those characterised by a combination of complexity (multiple variables interacting dynamically with a high degree of uncertainty),

[5]This point also holds for meta-analyses of RCTs, which frequently reach contradictory conclusions (Stegenga, 2011).

Table 2: The main steps of an Evidential Pluralism evaluation as set out by Shan and Williamson (2023, §12) and Trofimov and Williamson (2025).

	Task	*Key questions*
1.	Specify the causal claim.	What is the cause, effect, population and time-frame?
2.	Specify the correlation claim.	What are the potential confounders?
3.	Formulate specific mechanism hypotheses.	What are the key features of purported mechanisms of action, as well as counteracting and enhancing mechanisms?
4.	Search for and assess association studies.	What are the search terms? What are the inclusion criteria for the studies? How high quality is each individual study?
5.	Screen the need for a mechanistic evaluation.	What is the preliminary status of the correlation, general mechanistic and causal claims? Could evidence of specific mechanism hypotheses significantly change these preliminary determinations?
6.	Search for mechanistic studies.	Which features of specific mechanism hypotheses have already been established or ruled out? Which review questions should be used to find studies relevant to remaining features?
7.	Assess mechanistic studies.	How relevant are the population and variables of each study? How reliable are its methods? Does it implement these methods well? Are the results independently verified, consistent and robust?
8.	Assess specific mechanism hypotheses.	What is the status of each specific mechanism hypothesis in the light of the mechanistic studies?
9.	Assess the correlation claim.	Do specific mechanism hypotheses modify the preliminary status conferred on the correlation claim by association studies?
10.	Assess the general mechanistic claim.	Have alternative explanations of the correlation, such as bias and confounding, been ruled out? How well confirmed are the features of the mechanism complex? Can it account for the magnitude of the observed correlation?
11.	Assess the causal claim.	What is the minimum status of the correlation and mechanism claims?

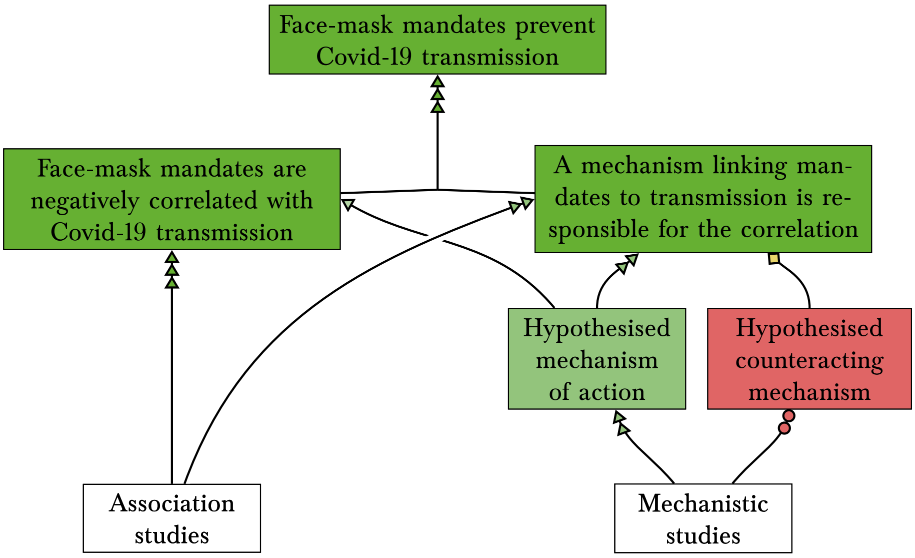

Figure 5: Evidential relationships in an evaluation of the effectiveness of face-mask mandates (Trofimov and Williamson, 2025, §6).

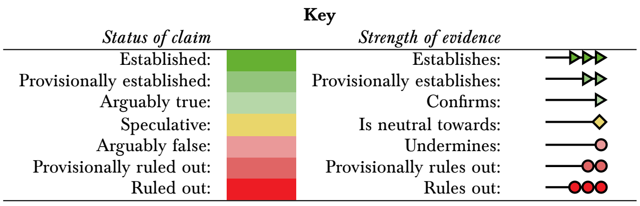

> urgency (decisions needed in days not years) and threat (the consequences of not acting could be catastrophic), mechanistic evidence has been mission-critical and RCTs difficult or impossible. Thousands of lives were likely lost as a result of what was incorrectly claimed to be an "evidence-based" approach—dismissing or downgrading mechanistic evidence, overvaluing findings from poorly designed or irrelevant RCTs, and advocating for inaction where RCT evidence was lacking. The pandemic is an epistemic opportunity for the EBM movement to come to better understand, debate and embrace EBM+.' (Greenhalgh et al., 2022, p. 253.)

One example concerns the evaluation of the effectiveness of mandates that required the wearing of face masks during the Covid-19 pandemic. The main orthodox review, which looked only at association studies, found the evidence to be inconclusive (Jefferson et al., 2023). A proof-of-concept Evidential Pluralism review, however, found that available association studies and mechanistic studies establish the effectiveness of face-mask mandates, as depicted in Fig. 5 (Trofimov and Williamson, 2025, §6). Following the normative evaluation guidance offered by Evidential Pluralism can thus make a substantial difference to resulting judgements of effectiveness.

Evidential Pluralism also provides a wider normative framework that can help researchers across the sciences determine how best to establish a causal claim of interest. Shan and Williamson (2023) argue that this has the potential to improve causal enquiry in the social sciences in general. Others have argued that Evidential Pluralism can help in more specific scientific contexts. For example, Iranzo and Pérez-González (2024) argue that Evidential Pluralism offers significant advances for decision making in public health, and Shaw et al. (2024) suggest that moving in the direction of EBM+ would improve the practice of thrombosis medicine.

Thus, a case can be made that Evidential Pluralism satisfies desideratum D3:

D3. Normativity. A science-friendly epistemological theory should offer the potential to improve scientific practice.

§4.4. Using concepts comprehensible to scientists

Evidential Pluralism appeals chiefly to the following concepts: establishing, evaluating, correlation, mechanism, causation, effectiveness, association study, mechanistic study. I shall give a flavour of their usage here; see Shan and Williamson (2023, §1) for more detail.

The term 'establish' is used in a familiar way: a proposition is deemed to be established when the standards are met for treating it as evidence, to be used to help assess further propositions. Establishing requires high confidence in a claim, as well as high confidence that the claim will not be undermined by future evidence. 'Evaluate' is also used in a familiar way: evaluating a claim is assessing its status in the light of the evidence: is it established or ruled out, or does it have some intermediate status?

'Correlation' is used in different ways across the sciences. Sometimes it is used in a narrow sense to refer to Pearson's correlation coefficient, which captures a kind of unconditional linear dependence; at other times, with much wider meanings. Evidential Pluralism uses 'correlation' in a very broad way to refer to any sort of difference between the probability of the putative effect when the cause is present and that in the case in which it is absent, when controlling for potential confounders. (In general, Evidential Pluralism tends towards broad construals of its key terms, rather than to technical formalisations that are to be found in some fields but not others.) 'Mechanism' is also understood in a broad way to encompass both mechanistic processes (spatiotemporally contiguous processes along which signals can be propagated) and complex-systems mechanisms (entities and activities organised in such a way that they are responsible for some phenomenon).

Evidential Pluralism is at pains to avoid giving the terms 'causation' and 'effectiveness' precise definitions that are grounded in metaphysical bias or methodological dogma. A probabilistic analysis of causation, in contrast, usually just defines causation to be some kind of difference between probabilities—i.e., as a correlation, in our sense. And in some sciences, an 'average causal effect' is simply defined as the average difference between outcomes, comparing the (counterfactual) situation in which all individuals are given the outcome and that in which the intervention is withheld from all individuals. Again, this can be construed as a kind of correlation. The worry is that these approaches attempt to circumvent the problems of causal enquiry by defining them away. Evidential Pluralism appeals instead to our standard informal notions of 'causation' and 'effectiveness'.

Evidential Pluralism uses 'association study' to refer to a study that seeks to estimate an association or correlation and 'mechanistic study' to refer to any study that sheds light on aspects of a mechanism. Again, these uses are familiar across the sciences, and a case can be made that Evidential Pluralism satisfies desideratum D4:

D4. Scientific Concepts. A science-friendly epistemological theory should appeal to concepts with which scientists are familiar and use them in the ways in which they are used in science, where possible.

There is, however, also something distinctive about the way in which Evidential Pluralism makes use of the terms 'association study' and 'mechanistic study'. As Fig. 1 shows, whether a study counts as an association study or a mechanistic study depends on the claim that is being evaluated, rather than the study design or the methods used to carry out the study. (For example, an association study for the claim that A is a cause of B counts as a mechanistic study for the claim that A is a cause of C, if B is hypothesised to be a mediating variable on a mechanism from A to C.) Thus the appeal to association studies and mechanistic studies is somewhat non-standard, in that it is more typical that studies are classified according to the methods they employ: quantitative versus qualitative studies, for example, or experimental versus observational studies. In part, it is by departing from these standard ways of classifying studies that Evidential Pluralism attempts to make progress on the problem of causal enquiry.

What Evidential Pluralism doesn't do is make things difficult for scientists by appealing to concepts in ways that are generally hard to grasp. According to certain externalist theories of evidence, for instance, your evidence base is unlikely to coincide with what you take to be your evidence base, and you are typically not in a position to identify the differences between the two (see, e.g., Williamson, 2000). These externalist theories of evidence can struggle to offer much practical advice to scientists, who can only go on what they take to be their evidence when trying to establish scientific claims. These externalist views of evidence thus depart from ordinary ways of thinking about enquiry to a much greater extent than Evidential Pluralism does (see Shan and Williamson, 2025, §3).

In sum, then, a case can be made that Evidential Pluralism has done enough to satisfy D4, though it will be a matter of debate as to whether any departures from standard ways of thinking about causal enquiry are compensated by the payoffs that the theory brings.

§4.5. Focus on what is important to science

As we have seen, the literature on Evidential Pluralism tends to draw on real examples from the health and social sciences, as well as from policy evaluation. This stands in marked contrast to much of the literature on causality, which tends to invoke toy examples in order to more easily test our intuitions. For instance, here is an example of Ned Hall's that favours a mechanistic view of causation over a counterfactual theory:

> Suzy and Billy, expert rock-throwers, are engaged in a competition to see who can shatter a target bottle first. They both pick up rocks and throw them at the bottle, but Suzy throws hers a split second before Billy. Consequently Suzy's rock gets there first, shattering the

> bottle. Since both throws are perfectly accurate, Billy's would have shattered the bottle if Suzy's had not occurred, so the shattering is overdetermined.
>
> Suzy's throw is a cause of the shattering, but Billy's is not. Indeed, every one of the events that constitute the trajectory of Suzy's rock on its way to the bottle is a cause of the shattering. But the shattering depends on none of these events, since had any of them not occurred the bottle would have shattered anyway, thanks to Billy's expert throw. (Hall, 2004, p. 235.)

Hall goes on to present similar examples that favour counterfactual theories over mechanistic theories. Indeed, it seems that almost all theories of causality are faced with problematic examples of this sort. While Evidential Pluralism seeks to set these problematic examples aside and focus on more typical and realistic scientific examples, the usual response to such counterexamples is to modify one's favoured theory of causality in order to somehow accommodate them.

This more usual approach has recently been roundly criticised by Timothy Williamson:

> I treat attempts to analyse knowledge, causation, meaning, and so on in more basic terms like Ptolemaic astronomy, to which new epicycles must continually be added to fix up anomalies. To invoke Imre Lakatos's useful category, many philosophers seemed unable to recognize a degenerating research programme when they saw one. (Williamson, 2024, p.x.)

Williamson takes many of these alleged counterexamples to be produced by heuristics that are less than perfectly reliable. Consequently, Williamson argues, the alleged counterexamples should not necessarily be taken at face value.

As discussed in §4.2, thinking in terms of correlations and thinking in terms of mechanisms can be viewed as heuristic devices for reasoning about causality, which, like other heuristics, are error-prone and can hence be construed as biases (see Fig. 4). Taken individually, each way of thinking can be construed as a fallible indicator of causation, and one shouldn't necessarily take its conclusions at face value. It is only when the two ways of thinking are integrated—as suggested by Evidential Pluralism—that causal judgements become more reliable.

There is some evidence from psychology for this heuristic view of correlational and mechanistic thinking. Thorstad and Wolff (2016) put forward a dual-process view of causal reasoning that sees correlational thinking as 'thinking fast', which is then moderated by mechanistic thinking:

> According to what we will call the Dual-Process Hypothesis of Causal Identification, identifying a cause involves two cognitive processes: 1) an automatic, intuitive process that identifies possible causes on the basis of perceptual cues (spatial and temporal) and 2) a slow, reflective process that identifies possible causes on the basis of causal inference, in particular, a consideration of possible mechanism. (Thorstad and Wolff, 2016, p. 919.)

They provide empirical evidence for this dual-process view, including a 'Jedi powers' experiment:

> In this study participants experienced an unexpected causal illusion on their way to our lab: a man appeared to open an elevator door [three times] by merely gesturing with his hands (Figure 1). Importantly, the man made no physical contact with the doors or with any of the buttons in the elevator. Unbeknownst to the participants, the doors were opened by a confederate outside of the elevator pushing the elevator button. Our main prediction was that the intuitive system would lead to feelings of causation that would ultimately be reduced by the reflective system. (Thorstad and Wolff, 2016, p. 920.)

This prediction was borne out by the results of the experiment.

Similarly, while Cheng (1993) and colleagues emphasise correlational thinking, Ahn et al. (1995) argue that:

> people do not treat the task of causal attribution as one of identifying a novel causal relationship between arbitrary factors by relying solely on covariation information. Rather, people attempt to seek out causal mechanisms in developing a causal explanation for a specific event. (Ahn et al., 1995, p. 299.)

On the other hand, mechanistic thinking on its own can itself be a kind of 'thinking fast'. Watching a moving billiard ball collide with a stationary ball is enough to infer causality, and Michotte (1946) generalised this observation by means of a range of animated 'launching experiments.' In each experiment, visual cues suggest a billiard-ball-like mechanism, and provide 'what is needed *to give the impression* of launching' (Michotte, 1946, p. 43). Michotte argued that this mechanistic thinking primes us to directly perceive causation in a single instance, as long as these visual cues are present. This can be viewed as a kind of thinking fast. If the experiment were repeated, any lack of a robust association might subsequently undermine this sort of fast causal judgement—a kind of thinking slow rather than fast.

Taking these results together, it is plausible that correlational and mechanistic thinking each provide a way of 'thinking fast' about causation; it is only when they are used in concert that we can be said to be 'thinking slow'. This view can be motivated by Evidential Pluralism, and it receives some further support in the psychological literature:

> beliefs about causal relations reflect at least two independent nonredundant sources of information: knowledge about the degree to which the cause and the effect covary and knowledge about causal mechanisms that mediate between cause and effect. These two sources of knowledge appear to be represented independently and appear to contribute differently to the subsequent evaluation of covariation-based data. Specifically, we found that covariation-based data were weighted more heavily for believable than for unbelievable candidates, but only when those candidates were believed to have a causal mechanism linking them to their effect and when the reasoners were making causal judgments, as opposed to estimating covariation. (Fugelsang and Thompson, 2003, p. 810.)

> the interpretation of covariation data is strongly influenced by mechanism knowledge. For example, learning about a covariation between

> a cause and effect has a stronger effect on the judged probability of a causal relationship when there is a plausible mechanism underlying the cause and effect (e.g., severed brake lines and a car accident) than when there is not (e.g., a flat tire and a car failing to start; Fugelsang & Thompson, 2000). ... These effects show that not only does mechanism information do something beyond covariation, but that it even constrains the way that covariation is used. (Johnson and Ahn, 2017, p. 129.)

According to this line of argument, then, we should not be taking Billy-and-Suzy-style examples at face value because they are the products of heuristics—of thinking fast. In science, we should be thinking slow, sacrificing speed in order to reach more accurate causal judgements. So a science-friendly theory of causal enquiry should be focussing on the examples that tend to matter to science and should not be too readily swayed by Billy-and-Suzy-style counterexamples:

D5. Typicality. A science-friendly epistemological theory should focus more on cases that are typical or important in science and less on far-fetched putative counterexamples.

This is exactly the strategy taken by advocates of Evidential Pluralism.

⋆

We have seen in this section that Evidential Pluralism can be viewed as a prototypical science-friendly epistemological theory, in virtue of its satisfying the five desiderata of §2. Perhaps the least clear-cut desideratum is D4: while Evidential Pluralism appeals to concepts with which scientists are familiar, it also appeals to distinctions (such as the distinction between association studies and mechanistic studies) with which scientists are less familiar. In its defence, however, any epistemological theory that leads to helpful new ways of thinking about enquiry will need to depart from the status quo in some way or another, and Evidential Pluralism is much less revisionary with respect to the conceptual core of science than many other theories.

§5
Conclusions

Desiderata D1–5 offer an account of what it is for an epistemological theory to be science-friendly. And Evidential Pluralism provides an example of a science-friendly epistemological theory, I have argued. If more of epistemology were science-friendly, the chasm that currently divides epistemology and the philosophy of science might be more easily bridged. As an example of such a bridge in the area of causal enquiry, Wilde (2021, 2024, 2025) has recently forged links between Evidential Pluralism and knowledge-first epistemology, a line of work that promises to increase the relevance of epistemology to philosophy of science.

If we accept the need for science-friendly epistemology, some interesting further questions arise. For example, are there other desiderata that a science-friendly epistemological theory should satisfy? Are there other examples of science-friendly epistemological theories?

Acknowledgements

I am grateful to Ken Aizawa, Rani Lill Anjum, Fraser Macbride, Thomas Uebel and Michael Wilde for very helpful comments. This research was supported by Leverhulme Trust grant RPG-2022-336 and UKRI CRCRM grant 25130.

Bibliography

Ahn, W., Kalish, C. W., Medin, D. L., and Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, 54(3):299–352.

Andersen, F., Anjum, R. L., and Rocca, E. (2019). Philosophical bias is the one bias that science cannot avoid. *eLife*, 8:e44929.

Aronson, J. K., Auker-Howlett, D., Ghiara, V., Kelly, M. P., and Williamson, J. (2021). The use of mechanistic reasoning in assessing coronavirus interventions. *Journal of Evaluation in Clinical Practice*, 27:684–693.

Benbrahim-Tallaa, L. (2019). Mechanistic data can play a pivotal role in IARC monographs evaluations when human data are less than sufficient. IARC Monographs publications. International Agency for Research on Cancer.

Berghofer, P. (2025). On the relationship between epistemology and science: synergies between experience-first epistemologies and agent-centered interpretations of quantum mechanics. *Synthese*, 205(1):45.

Berrington de González, A., Richardson, D., and Schubauer-Berigan, M., editors (2024). *Bias Assessment in Case–Control and Cohort Studies for Hazard Identification*. Statistical Methods in Cancer Research, Volume V. International Agency for Research on Cancer, Lyon. IARC Scientific Publication No. 171.

Bird, A. (2022). *Knowing Science*. Oxford University Press, Oxford.

Bradley, D. (2015). *A critical introduction to formal epistemology*. Bloomsbury, London.

Bueno, O. (2016). Epistemology and philosophy of science. In Humphreys, P., editor, *The Oxford Handbook of Philosophy of Science*, chapter 11, pages 233–251. Oxford University Press, Oxford.

Chakravartty, A. (2005). Causal realism: Events and processes. *Erkenntnis*, 63(1):7–31.

Cheng, P. W. (1993). Separating causal laws from casual facts: Pressing the limits of statistical relevance. In Medin, D. L., editor, *Psychology of Learning and Motivation*, volume 30, pages 215–264. Academic Press, San Diego, CA.

Clarke, B. (2011). *Causality in medicine with particular reference to the viral causation of cancers*. PhD thesis, Department of Science and Technology Studies, University College London, London.

Colaço, D., Buckwalter, W., Stich, S., and Machery, E. (2014). Epistemic intuitions in fake-barn thought experiments. *Episteme*, 11(2):199–212.

Collingwood, R. G. (1940). *An essay on metaphysics*. Oxford University Press, Oxford.

Craver, C., Tabery, J., and Illari, P. (2024). Mechanisms in science. In Zalta, E. N. and Nodelman, U., editors, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2024 edition.

Friedman, J. (2020). The epistemic and the zetetic. *The Philosophical Review*, 129(4):501–536.

Fugelsang, J. A. and Thompson, V. A. (2003). A dual-process model of belief and evidence interactions in causal reasoning. *Memory & Cognition*, 31(5):800–815.

Gillies, D. (2019). Should we distrust medical interventions? *Metascience*, 28(2):273–276.

Gillies, D. and Zheng, Y. (2001). Dynamic interactions with the philosophy of mathematics. *Theoria*, 16(3):437–459.

Greenhalgh, T., Fisman, D., Cane, D. J., Oliver, M., and Macintyre, C. R. (2022). Adapt or die: how the pandemic made the shift from EBM to EBM+ more urgent. *BMJ Evidence-Based Medicine*, 27:253–260.

Guyatt, G., Cairns, J., Churchill, D., Cook, D., Haynes, B., Hirsh, J., Irvine, J., Levine, M., Levine, M., Nishikawa, J., Sackett, D., Brill-Edwards, P., Gerstein, H., Gibson, J., Jaeschke, R., Kerigan, A., Neville, A., Panju, A., Detsky, A., Enkin, M., Frid, P., Gerrity, M., Laupacis, A., Lawrence, V., Menard, J., Moyer, V., Mulrow, C., Links, P., Oxman, A., Sinclair, J., and Tugwell, P. (1992). Evidence-based medicine: A new approach to teaching the practice of medicine. *JAMA*, 268(17):2420–2425.

Hall, N. (2004). Two concepts of causation. In Collins, J., Hall, N., and Paul, L., editors, *Causation and counterfactuals*, pages 225–276. MIT Press, Cambridge MA and London.

Howson, C. and Urbach, P. (1989). *Scientific reasoning: the Bayesian approach*. Open Court, Chicago IL, second (1993) edition.

IARC (2019). *IARC Monographs on the Evaluation of Carcinogenic Hazards to Humans: Preamble*. Lyon. https://monographs.iarc.fr/wp-content/uploads/2019/01/Preamble-2019.pdf.

Illari, P. M. (2011). Mechanistic evidence: Disambiguating the Russo-Williamson thesis. *International Studies in the Philosophy of Science*, 25(2):139–157.

Iranzo, V. and Pérez-González, S. (2024). Evidence and computer simulations in public health. *Global Philosophy*, 34(1):25.

Jefferson, T., Dooley, L., Ferroni, E., Al-Ansary, L. A., van Driel, M. L., Bawazeer, G. A., Jones, M. A., Hoffmann, T. C., Clark, J., Beller, E. M., Glasziou, P. P., and Conly, J. M. (2023). Physical interventions to interrupt or reduce the spread of respiratory viruses. *Cochrane Database of Systematic Reviews*, 2023(1):CD006207.

Johnson, S. G. B. and Ahn, W.-k. (2017). Causal mechanisms. In *The Oxford Handbook of Causal Reasoning*, chapter 8, pages 127–146. Oxford University Press, Oxford.

Lange, M. (2017). *Because without cause: non-causal explanations in science and mathematics*. Oxford University Press, Oxford.

Ludwig, D., Koskinen, I., Mncube, Z., Poliseli, L., and Reyes-Galindo, L., editors (2023). *Global Epistemologies and Philosophies of Science*. Routledge, Abingdon.

Maziarz, M. and Stencel, A. (2022). The failure of drug repurposing for COVID-19 as an effect of excessive hypothesis testing and weak mechanistic evidence. *History and Philosophy of the Life Sciences*, 44(4):47.

Michotte, A. (1946). *The perception of causality*. Basic Books, New York, 1963 edition.

Park, A., Steel, D., and Maine, E. (2023). Evidence-based medicine and mechanistic evidence: The case of the failed rollout of efavirenz in Zimbabwe. *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine*, 48(4):348–358.

Parkkinen, V.-P., Wallmann, C., Wilde, M., Clarke, B., Illari, P., Kelly, M. P., Norell, C., Russo, F., Shaw, B., and Williamson, J. (2018). *Evaluating evidence of mechanisms in medicine: principles and procedures*. Springer, Cham, Switzerland.

Popper, K. R. (1934). *The Logic of Scientific Discovery*. Routledge (1999), London. With new appendices of 1959.

Rubin, D. B. (2005). Causal inference using potential outcomes. *Journal of the*

American Statistical Association, 100(469):322–331.

Russo, F. and Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science*, 21(2):157–170.

Samet, J. M., Chiu, W. A., Cogliano, V., Jinot, J., Kriebel, D., Lunn, R. M., Beland, F. A., Bero, L., Browne, P., Fritschi, L., Kanno, J., Lachenmeier, D. W., Lan, Q., Lasfargues, G., Curieux, F. L., Peters, S., Shubat, P., Sone, H., White, M. C., Williamson, J., Yakubovskaya, M., Siemiatycki, J., White, P. A., Guyton, K. Z., Schubauer-Berigan, M. K., Hall, A. L., Grosse, Y., Bouvard, V., Benbrahim-Tallaa, L., Ghissassi, F. E., Lauby-Secretan, B., Armstrong, B., Saracci, R., Zavadil, J., Straif, K., and Wild, C. P. (2020). The IARC Monographs: Updated procedures for modern and transparent evidence synthesis in cancer hazard identification. *JNCI: Journal of the National Cancer Institute*, 112(1):1–8.

Schmitt, F. F. (2004). Epistemology and cognitive science. In Niiniluoto, I., Sintonen, M., and Woleński, J., editors, *Handbook of epistemology*, pages 841–903. Kluwer, Dordrecht.

Shan, Y. and Williamson, J. (2023). *Evidential Pluralism in the Social Sciences*. Routledge, Abingdon.

Shan, Y. and Williamson, J. (2025). Responses to criticisms. *Asian Journal of Philosophy*, 4(2):80.

Shaw, J. R., Nopp, S., Stavik, B., Youkhana, K., Michels, A. L., Kennes, S., Rak, J., and ten Cate, H. (2024). Thrombosis, translational medicine, and biomarker research: Moving the needle. *Journal of the American Heart Association*, 0(0):e038782.

Steel, D. (2008). *Across the Boundaries: Extrapolation in Biology and Social Science*. Oxford University Press, Oxford.

Stegenga, J. (2011). Is meta-analysis the platinum standard of evidence? *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 42(4):497–507.

Thorstad, R. and Wolff, P. (2016). What causal illusions might tell us about the identification of causes. In Papafragou, A., Grodner, D., Mirman, D., and Trueswell, J., editors, *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pages 1991–1996, Austin,TX. Cognitive Science Society.

Tooley, M. (1990). Causation: Reductionism versus realism. *Philosophy and Phenomenological Research*, 50:215–236.

Trofimov, A. and Williamson, J. (2025). Applying Evidential Pluralism to evidence-based law: EBL+. *Jurisprudence*, pages 1–44. Doi 10.1080/20403313.2025.

Weber, E. (2009). How probabilistic causation can account for the use of mechanistic evidence. *International Studies in the Philosophy of Science*, 23(3):277–295.

Weber, E. (2025). Evidential pluralism, epistemic causality and mixed methods research. *Asian Journal of Philosophy*, 4(1).

Wilde, M. (2021). Mechanistic reasoning and the problem of masking. *Synthese*, 199(3):6103–6118.

Wilde, M. (2024). A dilemma for the Russo-Williamson Thesis. *Erkenntnis*, 89(6):2437–2457.

Wilde, M. (2025). Evidential Pluralism and accounts of establishing. *Asian Journal of Philosophy*, 4(1):29.

Williamson, J. (2019a). Establishing causal claims in medicine. *International Studies in the Philosophy of Science*, 32(2):33–61.

Williamson, J. (2019b). Evidential Proximity, Independence, and the evaluation of carcinogenicity. *Journal of Evaluation in Clinical Practice*, 25(6):955–961.

Williamson, J. (2021). The feasibility and malleability of EBM+. *THEORIA. An International Journal for Theory, History and Foundations of Science*, 36(2):191–209.
Williamson, T. (2000). *Knowledge and its limits.* Oxford University Press, Oxford.
Williamson, T. (2024). *Overfitting and heuristics in philosophy.* Oxford University Press, Oxford.