

# Epistemic Interests and the Objectivity of Inquiry

Torsten Wilholt, Leibniz Universität Hannover, [torsten.wilholt@philos.uni-hannover.de](mailto:torsten.wilholt@philos.uni-hannover.de)

## Preprint

Article is forthcoming in *Studies in History and Philosophy of Science* (accepted for publication on Oct 8, 2021) and, once published, will be available at

<https://doi.org/10.1016/j.shpsa.2021.11.009>

© 2021. This manuscript version is made available under the CC-BY-NC-ND 4.0 license

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

## Abstract

This paper advocates for making epistemic interests a central object of philosophical analysis in epistemology and philosophy of science. It is argued that the importance of epistemic interests derives from their fundamental importance for the notion of objectivity. Epistemic interests are defined as individuated by a set of objectives, each of which represents a dimension of the search for truth. Among these dimensions, specificity, sensitivity, and productivity (the rate at which results are delivered per unit of time) are discussed in detail. It is argued that the relevance of productivity is often overlooked in debates about the ends and means of science. A definition of the objectivity of inquiry is proposed that takes the notion of epistemic interest as its starting point.

## 1. Introduction

My aim in this paper is to show that epistemology and philosophy of science ought to reinstate epistemic interests, which are understood and defined here in a specific sense that I shall seek to delineate, as a central subject of philosophical analysis. I maintain that the importance of the concept of epistemic interest is due to its intimate connection to a notion that was once crucial for epistemology, but that has come under criticism and has since been in crisis—the concept of objectivity. Understanding epistemic interests, I shall argue, will enable us to identify an important sense of objectivity that can provide an acceptable and much needed normative ideal and quality criterion for inquiry.

Objectivity—in any of the various senses in which the term has historically been used—is often an important concern in inquiry of whatever kind. In what follows, I will focus on scientific inquiry since it has traditionally stood as a model of objectivity and can

still provide particularly powerful examples which show that a perceived lack of objectivity continues to be an appropriate cause of alarm.

## 2. Objectivity in crisis

Objectivity is notoriously difficult to pin down as a concept. Attempts to capture it in philosophical analysis bring to light a dazzling variety of notions of objectivity, without revealing a clear 'brand essence' (Douglas, 2004; Megill, 1994).

On the one hand, there is an influential notion that reduces objectivity to so-called mind-independent truth. While common parlance, "mind-independent truth" is a slightly misleading expression, because what is at issue is not a certain kind of truth but rather the characteristic of being made true (or false) by a certain kind of fact. It may be true, for example, that the radiator seems hot to me, but this truth is not objective insofar as facts concerning the observer's state of mind figure among the relevant truth conditions. In contrast, the proposition that this radiator has a surface temperature of more than 60° centigrade can be construed as one whose truth conditions reside only in mind-independent facts. It is this sense of objectivity that underlies one of the central questions of modern Western philosophy, namely the question of how objective knowledge is possible in the first place. As such, this notion makes objectivity primarily a characteristic of propositions. Insofar as episodes of inquiry can be more or less conducive to knowledge of the relevant kind of propositions, they, too, can be qualified as more or less objective, but only in a derivative sense. Thus, the focus of this notion of objectivity is not on episodes of inquiry, but on kinds of propositions. Likewise, the only way in which objectivity in this sense is linked to *scientific* knowledge in particular is via the tenuous claim that (some) scientific methods are particularly apt for giving us access to mind-independent facts.

Objectivity is also associated with specific kinds of attempts to systematically weed out idiosyncratic influences on our system of knowledge. As a result, there is a notion of objectivity that can be described as a kind of invariance—in particular, invariance over change of perspective, but also invariance over change of system of reference and similar transformations.<sup>1</sup> Standardized modes of representation such as ground plans or cross sections of buildings and other technical artifacts, for example, embody the desire to emphasize objective information in this sense. Likewise, attempts to make maps that are less eurocentric strive for a faithful representation of ratios of areas or ratios of distances, and thus of magnitudes that, in contrast to proximity or preferred orientation, are invariant

---

<sup>1</sup> In his last book, Robert Nozick (2001) has elaborated a conception of objectivity for which this understanding is fundamental.

under different perspectives. Progress in methodology cannot eliminate the need for perspective or the necessity of a system of reference altogether, but it can serve to focus attention on those magnitudes and kinds of information that are invariant under change of perspective.

Relatedly, there is a sense of objectivity that signifies procedures that strive to minimize the risk of human error. Methods that serve to bypass or short circuit the human element in research are thus considered particularly objective (Daston & Galison, 2007, Chapter 3). Photography, for example, was soon after its invention hailed as a particularly objective method of producing scientific images.

Last but not least, another influential conception of objectivity is related to one specific kind of distortion that human agents are thought to bring into the process of inquiry—their values and interests. Often, objectivity used to be understood as value freedom.

It is not easy to make out a common core in this variety of uses. Synthesizing the constraints stipulated by each of them in an additive way to create a maximal notion of objectivity would lead us to some such conception as ‘whatever it is that is solely and uniquely determined by the object, without any influence of the subject’. Such a maximal conception does not seem to have much use except a negative one, to mark objectivity as something unattainable and illusionary—at least as far as science and human knowledge are concerned. But even weaker notions of objectivity have come under heavy criticism at least since post-war times.

An early criticism was a line of thought developed within critical theory, according to which in any kind of inquiry, human interests are required to give it purpose and direction. Bringing interests back into the epistemological picture was then thought to reveal that the most that traditional notions of objectivity could achieve was to define just one kind of epistemic interest among many that can legitimately be pursued.

Objectivity in the sense of a correct, subject-independent representation of fact had already come under criticism as a mere illusion in Max Horkheimer’s (1937/1999) portrayal of what he called “traditional theory”. This criticism was further sharpened in the 1960s, in the context of the positivism dispute. The critical theorists held that traditional epistemologists had, in their attempts to answer the question about the possibility of objective knowledge, neglected the ubiquity and permeating power of interests (Habermas 1968/1971). They had relied on a strict separation of theoretical from practical rationality, and thereby of the realm of the cognitive from the realm of human interests. However, this

perspective ignores the extent to which the very framework for every cognitive endeavor is constituted by the full phenomenon of human development (“Bildung”). Analyzing the “transcendental frame” in each case reveals the interests underlying it. Both Marx and the positivists, so the critical theorists believe, have only perpetuated these traditional mistakes by postulating a unity of science and thereby denying the decisive differences in epistemic interests underlying different kinds of inquiry. The early Habermas in particular argued that by identifying, distinguishing, and revealing the epistemic interests at work, critical theory could destroy the “objectivist appearance”, or “objectivist illusion” of science (1968/1971, p. 308).

If this criticism was ever seriously considered in mainstream philosophy of science, it was soon overshadowed by intellectual developments that began to unfold only a little later in the course of the onslaught of Thomas Kuhn, Paul Feyerabend, and other post-positivistic philosophers of science against the received view. Both were widely understood to have shown that experiment and observation could not be seen as neutral arbiters of truth in science and that this realization undermined the idea that science is based on objective fact. From the thesis that there is no paradigm-independent standpoint from which to judge the relative merits of competing hypotheses many popularizers of these new ideas inferred that values and interests had to play a role in theory choice.

This theoretical development has found its continuation in the more recent erosion of the value-free ideal. While Kuhn (1977) himself emphasized the relevance of a certain, limited set of values that were characteristic for science, later critics claimed that there was no principled distinction possible between epistemic and non-epistemic, or cognitive and non-cognitive values (Longino, 1990, 2002). Together with a second line of criticism that revived arguments from the 1950s and 1960s about inductive risks and the indispensability of value input for setting standards of evidence (Douglas, 2000, 2009), this resulted in a sweeping attack on value-freedom as an epistemic ideal. Today, the value-free ideal may be the one element from the range of traditional notions of objectivity that is most thoroughly dead, occasional desperate attempts at resuscitation notwithstanding.

Especially the problem that it seems impossible to manage inductive risks and set standards of evidence without recourse to non-epistemic values, as had been pointed out by Richard Rudner (1953) and forcefully elaborated by Heather Douglas, makes value freedom an ideal that seems very difficult to defend. An ideal of objectivity that appeals to value freedom could not possibly exclude extra-scientific values from all scientific practices, as deciding which questions to pursue as well as which results to apply to practical matters are obviously activities in which moral or political values may

legitimately be considered. The ideal would have to exclude such value-influences only from the kind of practices that are concerned with the testing and ultimately the acceptance or rejection of theories and hypotheses. The following would be a first approximation to such an ideal: ‘In objective research, moral, political, and other non-epistemic values must not exert any influence on the procedures that are decisive for the acceptance or rejection of hypotheses.’ But can questions regarding the acceptance of hypotheses be separated from questions regarding the application of results (where ethical and other values must be granted a legitimate role)? For every hypothesis in empirical science, however well confirmed it might be, there remains a certain probability of error. Therefore, whether a statement can be considered empirically established and acceptable depends not only on the evidence, but also on the standard against which its level of confirmation is measured. The sciences of statistics and inductive logic do not allow us to derive which standard is the right one to apply (that is which level of inductive risk can be tolerated) in a concrete situation. The only way to answer this question in a non-arbitrary manner is in light of the potential positive and negative effects that would come about as a result of the acceptance and dissemination of a true or false proposition, respectively. In Rudner’s (1953, 6) much-quoted words, “How sure we need to be before we accept a hypothesis will depend on how serious a mistake would be.” In this way, concerns about potential consequences and value judgments pertaining to them unfold a relevance for the assessment of validity claims.

It is one of Douglas’s (2000, 2009) important contributions to these arguments to have pointed out that the choices requiring evaluative judgment are not limited to a decision of whether or not to accept a hypothesis at the end of an episode of inquiry, but occur at all stages throughout the research process. More recently, Justin Biddle and Quill Kukla have noted that not all risks of error in epistemic practices can be calibrated in terms of how much evidence is sufficient to accept a hypothesis, as the framework of “inductive risk” seems to suggest (Biddle, 2016; Biddle & Kukla, 2017). Following them, I will use the term “epistemic risks” instead, to acknowledge that not all risks of error are always associated with some kind of propositional inference, and also to avoid implying a picture where evidence itself is assumed to be risk- and value-free while only evidence-based inference is risky and value-laden (on the latter point, see also Brown, 2013).

The developments outlined above have seriously undermined objectivity as an ideal and quality criterion for any type of inquiry. This is obvious for the conception of objectivity as value freedom, but it also pertains to other ideas associated with objectivity. The notion of objective truth as determined solely by mind-independent facts is not immediately affected by any of the criticisms I mentioned, but they do affect the possibility

to cash it out as an ideal that is applicable to episodes of inquiry rather than to propositions. It is no longer possible to straightforwardly turn the concept of objective truth into one of objective inquiry because, as the arguments surrounding epistemic risk and standards of confirmation illustrate, inquiry always has to strike a balance between striving for true belief and avoiding false belief. Both are indispensable ingredients in the pursuit of objective truth, each without the other is utterly pointless. In settling which level of epistemic risk to tolerate, we inevitably balance them against each other, and there is no neutral or logically correct way of doing so. This is where epistemic interests necessarily enter—a point that I will analyze more thoroughly below. For now, I would like to note that all this means that if you possess or defend a notion of objective truth, you do not automatically also have a conception of objective inquiry at your disposition. Equally, by leveling a convincing criticism against the objectivity of inquiry, you have not thereby also undermined the notion of objective truth (as the young Habermas might have thought).<sup>2</sup> Metaphysical realism is not the most salient issue with regard to the question of objective inquiry.<sup>3</sup>

Similar problems beset the strategy of describing objectivity as some kind of invariance. While we may try to single out certain *propositions* as expressing information that possesses such invariance, invariance of a *process* of inquiry seems neither possible nor desirable, except under technical transformations in the narrowest sense, such as between different units of measurement. In inquiry, a change of perspective or system of reference will often imply a different conceptualization of the consequences of getting it right and getting it wrong which will—and should—in turn have an effect on our standards for acceptable results. Similar considerations apply to the notion of objectivity as minimizing the risk of human error. As we have seen, there are always different kinds of possible errors. How we distribute our efforts at minimizing them over the different kinds is a fact that can typically not be explained or justified without reference to values and interests. More specifically, there seems to be no value-free *prima facie* justification for thinking that eliminating *human* sources of error should always be a priority.

---

<sup>2</sup> In his section on positivism, Habermas assumes that the ubiquity of interests also proves the correspondence theory of truth a hopelessly naive idea (1968/1971, p. 69).

<sup>3</sup> Which is not to say that it is no issue at all. If you deny mind-independent reality altogether, as some of the more radical constructionist anti-realists sometimes seem to do if you read them literally (e.g. Latour, 2000, p. 248), then there is not much merit in even starting debates about the objectivity of inquiry, or the trustworthiness of expertise.

### 3. Can we dispense with objectivity?

As a result of all this, there seems to be no acceptable conception of the objectivity of inquiry in sight for present day epistemology and philosophy of science. Should we perhaps conclude: All the better, we have finally shed the “objectivist illusion”, and good riddance? This was supposed to be a liberating moment in the history of thought, but the lacuna that an abandonment of objectivity as an epistemological ideal and quality criterion leaves behind proves debilitating for epistemology as a critical intellectual enterprise.

In particular, it would mean that many cases of biased, distorting research could no longer be criticized as epistemologically problematic, but only as problematic with regard to non-epistemic criteria. Without a substantial ideal of objectivity (as an *epistemic* ideal) researchers could not be criticized for allowing idiosyncratic value influences (such as the profit interests of a particular branch of industry) to have an impact on their research, but only for permitting the ‘wrong’ kinds of values (i.e., others than those preferred by the critic).

Consider again the observation that there is no value-free way of determining the level of epistemic risk that can and should be tolerated when accepting the result of a particular line of inquiry. A decision to apply a particularly strict standard for a (positive) result will automatically imply a looser standard for negative results or an increased risk for the inquiry to end up without any result at all. Typically, each and every methodological decision offers an opportunity for trading these different kinds of risks off against each other and opens up a whole continuum of different ways in which the balance between them could be settled. Must each and every possible distribution of epistemic risks be regarded as epistemically equally acceptable? Should differences between them only be debatable as differences of opinion with regard to the ethical evaluation of potential consequences of error?

Consider, for example, the choice of animal model in a toxicological trial. Each type of laboratory animal has different characteristics that are potentially associated with different probabilities of detecting health effects of a given substance of concern. Industry scientists might, for example, choose a type of animal that is known for its specific robustness and insensitivity with regard to effects caused by the class of substances at issue in order to decrease the probability of a positive result.<sup>4</sup> If challenged, they might cite concerns about false positives and possible over-regulation as a justification of their choice.

---

<sup>4</sup> This hypothetical example is modeled on a concrete episode described in Wilholt (2009).

If there was no such thing as objective inquiry, then the only way of criticizing the actions of the industry scientists in this scenario would be to take the debate to the level of value judgments—for example by demanding that the avoidance of false negatives should have a higher priority in toxicological research because of their potentially grave consequences for consumers. However, there seem to be instances along the lines of the case imagined here where the manipulation of the distribution of epistemic risks can no longer be justified with reference to differing opinions about the seriousness of mistakes, but where instead the resulting error probabilities amount to a shortcoming, the acknowledgment of which does not require one to possess a particular value outlook—because it is an epistemic rather than a moral failure. It undermines the epistemic trustworthiness of the respective episode of inquiry. My aim is to articulate a conception of objectivity that will enable us to describe the particular sense in which such examples are cases of epistemic shortcoming and that, at the same time, does justice to what is convincing about the criticisms of traditional notions of objectivity.

#### **4. Epistemic trust and the relevance of shared standards and practices**

In undertaking to overcome the impasse that these criticisms have left us with, I take my cue from the observation that a problem of unmistakably epistemic character emerges when idiosyncrasies of an episode of inquiry undermine its *trustworthiness* as an information-providing process.<sup>5</sup> The notion of trust invokes the social dimension of inquiry. All significant cognitive enterprises probably involve some level of cooperation or at least belief based on the testimony of others.<sup>6</sup> This is certainly true of the pursuit of knowledge in the sciences (Hardwig, 1991). If trust among different actors comes into play, this should place some constraints on their individual, value-based choices that impinge on the distribution of epistemic risks. Inquirers who know that they gauge epistemic risks similarly will be able to invest more trust in each other's results. For placing at least some trust in the results of researchers who are implementing a different value-based assessment of tolerable epistemic risks than I do, I will at the very least have to be able to know *how* they distribute epistemic risks.

---

<sup>5</sup> The centrality of the concept of trust for understanding scientific objectivity has been emphasized by Fine (1998, pp. 17-19) and Douglas (2009, p. 116).

<sup>6</sup> Coady (1992) forcefully argued for the indispensability of testimony for the justification of even our most mundane beliefs.



All this may be seen as creating a presumption in favor of shared standards for the distribution of epistemic risks, justified in virtue of the necessity of mutual trust.<sup>7</sup> However, it would be a mistake to view the objectivity of inquiry as defined by the shared stance of a community of inquirers. Trying to understand the balance of risks of error, such as it is struck by the majority of researchers in their methodological practices and their conventional standards, as the solely binding standard and only deviations from it as bias and distortion, would mean drawing too simple an image of the interplay between values and knowledge communities. This can best be shown using examples of the kind that I will call “pervasive distortions”. Under this label, I want to discuss cases in which entire communities of researchers, rather than just individuals or factions within a community, collectively adjudicate epistemic risk differently than they ought to.

Examples for pervasive distortions are more difficult to identify and discuss for philosophers of science than examples of particular biased studies, because we cannot rely on internal disciplinary controversy as an indicator that something is going wrong. The clearest examples are perhaps those that we can judge with the benefit of hindsight, such as the case of androcentric bias throughout most of the history of the life sciences. To pick a more specific aspect of this bias, consider the fact that male-dominated medical sciences have over centuries hardly skipped an opportunity to allegedly demonstrate by means of physiological facts that women were unsuitable for intellectually demanding tasks: Women’s brains were too cold and soft, the female skull too small for a powerful brain, intensive use of the brain adversely affected the female reproductive capacity, lack of lateralization made the female brain inferior for specific tasks - whenever one hypothesis proved scientifically untenable, the next approach stubbornly stepped onto the scene (Kourany, 2016, pp. 780-782).

Suppose for the sake of argument that this history is not grounded in a complete disregard for the truth on the part of the researchers. Then, a likely explanation is that it is due to a consistently biased management of epistemic risks. Researchers accepted far higher risks of false positives with regard to hypotheses concerning the alleged inferiority of women than they ought to have done, and thus jumped to conclusions on the slightest shred of putative evidence.

---

<sup>7</sup> I myself have used this idea to define a particular type of bias in terms of an infringement of a research community’s conventional standards (Wilholt, 2009, 2013, 2016); the suggestion has been taken up by John (2015) and Biddle and Leuschner (2015). In what follows, I do not intend to criticize this as a viable approach for identifying *some* problematic kinds of distorting bias, but only to point out that it cannot be straightforwardly generalized to yield a useful characterization of objectivity.

Note, however, that the normative judgment on what risk the researchers ought to have found acceptable in this case cannot be grounded in the implicit or explicit conventions of the respective research community. Given the consistent androcentrism of medical science throughout much of its history, the problematic epistemic risk management may not have even been in conflict with the implicit consensus of the research community as it was inscribed in its methodological practices.

Another example of the phenomenon of pervasive distortion is conservative bias in climate science. Several independent analyses of past assessment reports of the Intergovernmental Panel on Climate Change (IPCC) have concluded that, compared to observational data, predictions released with the reports have tended to underestimate some key parameters of global warming (Rahmstorf, 2007; Committee on Strategic Advice on the U.S. Climate Change Science Program, 2009; Allison et al., 2010; Freudenburg & Muselli, 2010; Brysse et al., 2013). Notably, this included predictions concerning sea-level rise, but also those on CO<sub>2</sub> emissions and atmospheric concentrations, on the melting of continental ice-sheets and sea-ice decline. These underestimations of the IPCC reports seem to reflect a conservative tendency that is already present in the research literature. It has been suggested that this is an effect of the activism and media presence of climate change contrarians, whose criticism often takes the form of personal attacks aimed at undermining the credibility of climate scientists. As a result of this, climate scientists often find themselves in a defensive position. Scientists go out of their way to avoid the impression that they ignore criticism, and thus ideas concocted by contrarians from outside the discipline (such as that of a purported post-1998 ‘hiatus’ of global mean surface temperature rise) have demonstrably seeped into serious scientific discourse (Lewandowsky et al., 2015). Apparently, many scientists also take special care to steer away from false positives in their predictive conclusions, lest they be branded as ideologically biased ‘alarmists’—they prefer “erring on the side of least drama”, as Brynne et al. (2013) have put it. From a philosophical point of view, this kind of pervasive distortion is interesting because it is not rooted in the climate scientists’ own prejudice. The skewed balance of epistemic risks does not reflect the scientists’ own value judgments, but rather those of their adversaries. Nonetheless, assuming that the conservative bias is already ingrained in the practices of mainstream climate scientists, any lack of objectivity that lies therein can obviously not be diagnosed by referring to the standards of the respective research community.

Lest it appear that I am presupposing that pervasive distortion is generally caused by the preferences of some actors involved for certain types of research results, I would like to

briefly refer to another example. A particular kind of methodological bias has drawn much attention lately: In experimental fields (especially those with a high prevalence of underpowered studies), incentives and conventions favoring positive results could lead to an overrepresentation of positive findings in the published research literature. This has been related to the phenomenon of failing replications as it is being discussed under the label “replicability crisis” (Romero, 2019). The present social organization of research seems to permit and incentivize scientists to exploit their “researcher degrees of freedom” (and, arguably, to engage in “questionable research practices”) in ways that are likely to contribute to a high false positive rate (Simmons et al., 2011; John et al., 2012; Ioannidis et al., 2014). It is debatable, for a variety of reasons, whether it is desirable or even possible, at least in certain areas of cutting-edge experimental research, to achieve a substantially improved reproducibility of results or substantially lower false-positive rates (Bird, in press; Feest, 2019). But the mere fact that a large number of researchers agree that problems with the reproducibility of results add up to a “crisis” (Baker, 2016) shows that the conventions and standards of the respective disciplines do not manage to rule out false positives to the extent that scientists themselves expect them to—or used to expect them to, prior to becoming aware of the problems.

When an individual study, by virtue of its methodological choices, tacitly deviates from the way in which error risks are routinely managed within the relevant research community (as required by explicit standards, or simply because it is consistent and unquestioned practice), then this can inherently undermine epistemic trust. However, the examples of pervasive distortions discussed in this section all show that this is not the only relevant type of threat to trustworthiness and objectivity, and that there are cases in which the established standards and practices themselves come under suspicion of being the reason for a lack of objectivity. These require an analysis that does not rely on shared standards and practices alone as the basis of trustworthiness. I will argue that reviving the concept of epistemic interest (in a new, more precise sense) will be of great help in approaching such an analysis.

## **5. No objectivity without epistemic interest**

When an entire research community shares biased practices and standards, trust among its members may continue to work perfectly well, but its results are arguably no longer trustworthy from without. This seems to suggest that, over and above following the *same* practices and standards, in order to merit epistemic trust, communities in some sense need to follow the *right* ones. How might the relevant sense of rightness be grounded?

From a traditional perspective, it might seem attractive to try and provide an answer to this question by grounding the objectivity of an episode of inquiry in the notion of objective truth: The right standards and practices are conducive to the attainment of (mind-independent, invariant, value-free) truth. However, the cases we just discussed may serve again to illustrate the problem with this approach that we identified in section 2.

To introduce a bit of convenient jargon, gleaned from the terminology of medical testing, let us call the reliability of positive results (i.e., the probability that a positive result is true) within a particular study or research methodology its *specificity*. Correspondingly, let us use *sensitivity* for the reliability of negative results within a given research methodology. It is important to note that these two central dimensions do not yet completely describe the degrees of freedom of epistemic risk management. A third dimension along which different projects of inquiry into the same question can vary is their *productivity*—the rate at which they produce results of any kind at all, irrespective of their positive or negative character, over a certain time span.

To emphasize the importance of the productivity dimension (often overlooked in discussions about epistemic risks), take the example of experimental psychology, one of the fields of concern in the replicability crisis. To claim that published results in this field contain too many false positives is to claim that experimental research in psychology is (or has been) lacking in specificity. However, this does not imply that low specificity must have been accepted in exchange for high sensitivity. In fact, an analysis of published non-significant results in psychology has indicated that many of them may be false negatives (Hartgerink et al., 2017). There is at present no reason to believe that experimental research in psychology is displaying a high degree of sensitivity. Instead, it seems that both specificity and sensitivity have been traded off for an increased productivity of psychological research.

Would this mean that psychological research is insufficiently geared towards the truth? No, because productivity is a dimension of the search for truth just as sensitivity and specificity are. It is almost trivial to achieve high sensitivity and specificity if you are prepared to sacrifice productivity—for example by adopting a research strategy that only deals with obvious cases and suspends judgment on all the rest. Aiming for the truth implies aiming at reliable results, but it also implies aiming at results. A research program that would place no value at all on productivity could not convincingly claim to be concerned with the pursuit of truth. However, the fact that many psychologists themselves are alarmed about the reported replicability problems indicates that existing practices have been *overemphasizing* productivity at the cost of specificity and sensitivity in a way that

crosses a line. It is not a lack of truth-directedness that seems to be a problem, but a somehow inadequate balance between the different dimensions of pursuing the truth.

Similarly, diagnosing conservative bias in present-day climate science does not imply a charge that researchers are not fully committed to truth-conducive inquiry. It merely means that existing practices place an exaggerated priority on specificity—this time to the considerable detriment of sensitivity.

Some of the 20<sup>th</sup> century criticisms of objectivity came to the conclusion that truth is too strong a notion; that defining the aims of science in terms of truth would mean to make them unattainable by definition. The above considerations show that the opposite is the case. Truth alone is too weak a notion for holding a substantial and helpful conception of the objectivity of inquiry together.

It is for this reason that redefinition of the concept of epistemic interest is needed. The definition must explicitly refer to the orientation of inquiry towards truth but at the same time emphasize that within the confines of this orientation, different specific kinds of directedness exist. Each such kind of directedness represents a specific epistemic interest. An epistemic interest in this sense is a set of objectives characterizing the exact manner in which a given activity is directed at finding the truth. More precisely, this set of objectives must include: a question (or set of questions) that the investigation is devoted to finding the true answer to, a target specificity, a target sensitivity, and a target productivity.

For some contexts, it might be useful to individuate epistemic interests further, as there is another dimension to the pursuit of truth that is not yet reflected in the above set of objectives. When science pursues a question, finding a full, maximally informative answer is often not possible (at least within predictable time frames). Instead, research typically delivers (and strives to deliver) partial answers. Partial answers to the same question may differ along several dimensions such as comprehensiveness and precision. For illustration, consider the following answers to the question “Where in the world does *Rhinoceros sondaicus* live in the wild?”—“In Asia.” “Only in Indonesia.” “Some can be found on the island of Java.” “All remaining individuals live on the westernmost tip of Java, in Ujung Kulon National Park.” Epistemic interests in the same question and aiming for the same specificity, sensitivity, and productivity may still differ in the precision and comprehensiveness of the partial answers that they strive for. For the remainder of this paper, I will nevertheless abstract from this aspect and treat epistemic interests as individuated on the basis of question pursued, target specificity, target sensitivity and target productivity.

Restricting the characteristics that individuate epistemic interests to such traits as can be understood as dimensions of the search for truth is intended to accomplish nothing more than to do justice to the truism that “in pursuing the answer to a question, we are automatically pursuing the truth on that question” (Sosa, 2003, p. 157). In particular, this restriction does not amount to a commitment to the claim that truth is ‘the aim of science’.<sup>8</sup> In particular, it in no way entails that if  $p$  is an unknown truth, some scientist (or ‘science’) must develop an epistemic interest concerning  $p$ . In effect, the restriction stipulates that *if* any cognitive subject pursues an interest that concerns a particular question, it only counts as an *epistemic* interest if certain minimal evaluative attitudes are reflected in the actions that constitute its pursuit—specifically, preferences (however slight) for getting true results over getting no results and for getting no results over getting false results.<sup>9</sup>

For such preferences to be meaningful at all, the truth or falsity of the outcome of an investigation must be considered constituted by something that is independent of the course of the investigation itself and its outcome. In this sense, the conception of epistemic interests proposed here presupposes alethic realism. However, it is by no means committed to scientific realism, but is fully compatible, for example, with anti-realist positions such as constructive empiricism (van Fraassen, 1980). The latter would merely amount to a restriction on the epistemic interests pursued in the sciences, confining them in the first dimension (the set of questions pursued) to questions about observable objects.<sup>10</sup>

On the basis of this revised notion of epistemic interest, we can now make a new attempt at understanding objectivity as a property of episodes of inquiry—namely, as a form of match between the method of investigation and the relevant epistemic interest. A good match in this respect is what underlies the trustworthiness of science as an epistemic enterprise. This would mean that the ubiquity and diversity of epistemic interests is not

---

<sup>8</sup> For some deserved criticism of talk about the ‘aim of science’, see Rowbottom, 2014.

<sup>9</sup> More precisely, a minimal commitment to truth requires preferring, in any world in which  $p$  holds, that the inquiry come to the result  $p$  over an inquiry with no result, and in turn preferring that over the result not- $p$ , and conversely, in any world in which not- $p$  holds, preferring an inquiry with result not- $p$  over an inquiry with no result, and preferring that over the result  $p$ . No particular kinds of preferences between the worlds realizing  $p$  and non- $p$  are required, because these have no meaning for methodological decisions (insofar as these are rationally made in the light of preferences). For a Bayesian model that justifies this last point and shows that the weak preferences mentioned above are sufficient for methodological decisions made in their light to be expected to lead to an investigation with non-vanishing specificity and sensitivity, see Wilholt (2013).

<sup>10</sup> Since van Fraassen (1980, pp. 12-13, cf. p. 82) also acknowledges a pragmatic dimension of theory acceptance, meaning that scientists who accept a theory use its full conceptual repertoire (not only what it says about observables) to confront future phenomena, seek explanations, and so on, constructive empiricists adopting the language of epistemic interests might say that scientists also in some sense ‘pursue’ questions about unobservable things, but that this form of pursuit need not constitute genuine *epistemic* interests (and *will* not amount to it for scientists who are appropriately modest about the achievable epistemic objectives of science).

undermining the very possibility of scientific objectivity as critical theory once suspected. Rightly understood, objectivity is a relation of appropriateness between a collective effort of inquiry and its underlying epistemic interest.

But what is the ‘relevant’ or ‘underlying’ epistemic interest that scientific research needs to reflect in order to count as objective? A straightforward way to answer this question would be to stipulate that researchers ought to adjust their methodologies to the epistemic interests prevailing in society.<sup>11</sup> However, this might easily lead to counter-intuitive consequences. While there is good reason to demand that researchers take the public’s views on risks and potential benefits *into account*, they should not just emulate the public’s views. Consider the example of a 2007 survey which showed that the public is considerably *less* concerned about the potential new health risks arising from research on nano-materials than the scientists involved in the pertinent research (Scheufele et al., 2007). It would not be a convincing proposal to suggest that researchers studying potential health risks of nano-materials ought to lower their sensitivity targets in accordance with the more relaxed attitudes of the public.

Instead of requiring that researchers ought to adjust their methodology to *prevailing* interests, one thing that can and must be expected of them is to *ceteris paribus* conform to the epistemic interests that are *characteristic for the kind of inquiry that they are publicly perceived to be engaged in*. An episode of research is objective exactly to the extent that it is effectively designed in every detail and in every respect to serve precisely these epistemic interests. This is a standard against which individual research teams as well as entire research communities (together with their own methodological conventions) can be measured.

Conservative bias in climate science is a phenomenon of concern because the epistemic interest that climate science is perceived to be serving would demand that it avoids a systematic underestimation of effects. It is not that there is not *some* epistemic interest that could be served by an undertaking that gives all the priority to avoiding false positives at the risk of systematically underestimating future developments of global warming – but that is not the epistemic interest that almost everybody, including climate scientists themselves, associate with the collective enterprise of climate science.

Objectivity is a normative concept. The normative force of objectivity, understood along the lines that I am suggesting, can be explained with reference to the indispensability

---

<sup>11</sup> In a similar spirit, John (2019) has demanded that the justification of published findings be compatible with the values of the putative audience, and Intemann (2015) suggests that influence of non-epistemic values on research is legitimate if it promotes “democratically endorsed” aims.

of epistemic trust. It is not the case that objectivity in this sense guarantees that everyone can trust the results of objective research without reservation. After all, the epistemic interests of the individual may well diverge from the epistemic interests that the respective research is serving according to public perception. But objective inquiry in this sense enables users of information to reliably *assess* the degree to which results can be expected to fit their own epistemic interests.

The epistemic interest that an episode of inquiry is ‘perceived to be serving’ is clearly a more or less vague concept. That is not to say that it may not even be quite precisely defined for some cases, especially when the type of inquiry is highly formalized, when it has been well established as a source of public information for an extended period of time or has an extended and transparent track record. But for some cases the perceived epistemic purpose of a branch of scientific research may be so vague as to correspond to a rather wide scattering of possible epistemic interests. (Note that even if the perceived epistemic interest is really a rather wide scattering of epistemic interests, it does not mean that there are not epistemic interests that are definitely and recognizably not included. Vagueness does not make a concept useless; cf. van Fraassen, 1980, p. 16.) The concept of objectivity thus inherits this vagueness by way of the definition that I propose. I think this is how it ought to be, as objectivity is in fact more difficult to assess when the exact epistemic purposes of an investigation are shrouded in imprecision. In extreme cases, it may become impossible to achieve.

In a similar vein, the public may be so deeply divided with regard to what they perceive as the epistemic interest that a given investigation is serving that one cannot reconstruct a shared conception, however vague, of its perceived epistemic interest. I believe this kind of disagreement is less common than conflict over the epistemic interest that scientific research *ought* to be serving, which is a different matter and does not affect my definition. But it is certainly possible, and in that case, too, objectivity becomes unattainable. Again, I think this is appropriate: In whichever way one defines objectivity, there will always be conceivable constellations of deep disagreement under which meaningful exchange about the objectivity of inquiry breaks down.

Both observations entail a normative corollary: To achieve objectivity, epistemic interests should ideally be transparently communicated. For this purpose, the researchers themselves must first make a conscious determination of their epistemic interests. This is related to another normative ramification: Objectivity requires that the design and course of the investigation is guided by the epistemic interests—and not vice versa. Confusing this order and adapting the ostensibly pursued epistemic interests in the light of obstacles or



unforeseen events of the research process is a source of many typical questionable research practices.

## 6. Objections and rejoinders

*Objection: An episode of inquiry that follows one particular epistemic interest represents only one perspective. An adequate definition of objectivity must require multiple points of view; it is better to define objectivity as pluralism of approaches and value influences.*

Especially Helen Longino (1990, 2002) has pioneered a strategy in which the concept of objectivity is in effect abandoned at the level of individual researchers and then regained at the level of research communities—in the form of an ideal of diversity and plurality. A first rejoinder to this objection is to point out that the pluralist strategy of defining objectivity does not address the challenge that I have outlined in section 3. A conception of objectivity that resides on the level of the community is ineffective when it comes to the goal of finding a basis on which *individual* scientists or groups can be criticized on epistemological grounds, e.g., for deliberately choosing a research design or a publication strategy that strongly favors positive results over negative ones.

A second reply is independent of the particular philosophical goals that I have prioritized in this paper: While science thrives on controversy, scientific communities do not, as a general rule, remain divided with respect to important questions of their disciplines over indefinite periods of time. At least sometimes they converge on a consensus. What do we make of this in the context of the pluralist framework? If objectivity-as-pluralism is an ideal, why ever strive for consensus? (Cf. Carrier, 2013.) Does settling a dispute go against the ideal of objective inquiry? This question may seem lame, but it is very difficult to answer if objectivity is *defined* in terms of pluralism. The difficulty suggests that pluralism does not work as a definition of objectivity. More precisely: While the pluralist ideal may correctly identify important instrumental preconditions for scientific objectivity, it fails to give a satisfactory answer to the question of what objectivity, as a desirable characteristic of research processes, consists of. To see this more clearly, we may pose the question: Why is objectivity a desirable characteristic of a collective epistemic enterprise? If objectivity is defined as pluralism, the best answer might be: in order to facilitate mutual criticism. But why do we want mutual criticism? Pluralism and mutual criticism are clearly features of the scientific enterprise that are not desired for their own sake, but because they facilitate something. Only if and when we can say what that something is have we found the substantial notion of objectivity that we are looking for.

*Objection: Why so complicated? Science can secure the trustworthiness of its results by following strict standards. Strict standards make sure that individuals have good reason to trust, no matter what their personal epistemic interests are.* Both Carrier (2010) and John (2015) have argued along these lines, where strict standards are understood to mean a painstaking avoidance of false positives—in the terminology used earlier: a high specificity. Both Carrier and John work from the idea that a high degree of specificity guarantees that results of the research will be trustworthy for users who, in my terminology, share epistemic interests with lower specificity targets. High standards would then be downward compatible, in a manner of speaking.

The central limitation of this argument is that it only holds for positive results. For negative results, it is sensitivity rather than specificity that matters for compatibility. Also, the argument overlooks that users have productivity targets in addition to their specificity and sensitivity targets. In other words, they may be in urgent need for information, a need that is badly served by deferring judgment until an extremely high degree of specificity can be guaranteed. And given that science never operates with unlimited resources, higher specificity means either lower sensitivity or lower productivity, or both. The case of conservative bias in climate science illustrates the point. Arguably, conservative bias means that specificity targets are excessive, meaning that much needed information is not made available, or is made available too late. This makes the totality of the output of scientific research on this subject less trustworthy for an audience with higher productivity targets, and can potentially negatively affect their actual trust.

Stipulating that high specificity best serves the interests of science is not an interest-free proviso by which science stays on the safe side. It is itself the expression of a particular kind of epistemic interest. In methodology, there is no such thing as the safe side.

Another good case in point are the many problems that have been identified with the methodological ideals of evidence-based medicine (EBM). One way to pinpoint a core deficiency in its hierarchical methodology that places large randomized controlled trials at the top is that it gives almost all the priority to specificity, much less to sensitivity, and none at all to productivity. Doctors and patients alike do need reliable positive results from medical research, and EBM places a high premium on this need. But the need for reliable results needs to be balanced against the need for results. In this regard, EBM may fail the prevailing epistemic interests of the public in many cases (Worrall, 2002; Goldenberg, 2006; Borgerson, 2009).

*Objection: When epistemic interests are individuated by objectives that are characterized as 'dimensions of the search for truth', then what are the references to critical theory good for? This has nothing to do with epistemic interests in their sense anymore.* From my point of view, the proposed account is, on the contrary, also an attempt to revive certain core Habermasian ideas: For every kind of inquiry, there is an underlying epistemic interest that is constitutive for the knowledge-generating process. It is wrong to conceive of these interests as distorting or interfering with an imaginary ideal of inquiry. There is no coherent and non-arbitrary conception of inquiry that does without epistemic interests (cf. Habermas, 1968/1971, p. 210).

Admittedly, there are also striking differences: Habermas meant to identify three types of inquiry and correlate them with three epistemic interests (“empirical-analytical” inquiry corresponds to “technical” interest, “historical-hermeneutical” inquiry to “practical” interest and “critically oriented” inquiry to “emancipatory” interest). In contrast, my proposed redefinition implies a multi-dimensional continuum of epistemic interests. Perhaps more crucially, Habermas (1968/1971, pp. 69, 307) also thought that the notion of objective truth was thereby discredited. He thought that the idea of an ideal of inquiry that is free of values and interests is a straightforward consequence of a realist notion of truth, and that the latter must therefore be abandoned with the former. But since the claimed entailment does not hold, the path is open for a truth-based conception of epistemic interests.

## **7. Conclusion**

I have argued that, properly understood, objectivity designates the proper relation of an episode of inquiry not to an object, but to an objective (or set of objectives). To my mind, this is more than just a pun, or a piece of revisionist etymology. What is important is that the proposed definition captures the insight that objectivity is a kind of fit—one that has so far tended to be overlooked. The recent philosophical debate about values in science has largely been framed as one about how the ways in which value judgments involved in research are often non-transparent and difficult to account for. This makes it easy to overlook that even if we were able to precisely account for all the value influences and precisely assess the epistemic risks involved, results would still *be* value-laden—and therefore, depending on the epistemic interest that the research in question is supposed to be serving, more or less useful and trustworthy. Knowing exactly how an episode of research deviates from an intended epistemic interest is crucial for *calibrating* the appropriate level of trust, but it does not make it more capable of *servicing* that particular

epistemic interest (or of justifying a *high* level of trust from people who share it). For that, what is ideal is as close a match as possible.

I have proposed a relativized notion of objectivity—but it is appropriately relativized. Some might balk at a notion of scientific objectivity that makes it relative to epistemic interest. But note that an epistemic interest is not just any kind of interest. It is a set of objectives that all represent dimensions of the search for truth. All this is compatible with, in fact premised on, a realist conception of truth. There is no danger that the definition will make talk about objectivity arbitrary or unduly permissive.

It is also a conception of objectivity that takes episodes of inquiry as the proper category to which the notion is first and foremost applied. Finally, a good match between methodology and publicly identifiable epistemic interest underlies the trustworthiness of science as an epistemic enterprise, whereas a lack of such a match undermines trust by making the calibration of appropriate levels of trust impossible. The proposed notion thus captures the intimate relation between objectivity and trustworthiness.

## **Acknowledgements**

I am grateful for having had the opportunity to present predecessor versions of bits and pieces of this paper at the following workshops and conferences, and in each case thank participants for helpful feedback: *Biases in Science* (Munich Center for Mathematical Philosophy, 2019), *Scientific Authority & Scientific Fraud* (Universitat Autònoma de Barcelona, 2016), *SPSP 2016* (Glassboro, NJ, 2016), *GWP2016* (Düsseldorf, 2016), *Values and Theories* (ZiF, Bielefeld, 2015). I would also like to express my gratitude for helpful comments and suggestions I received during the review process for this journal. Leonie Wiemeyer proofread the paper for linguistic quality; she deserves thanks for numerous improvements.

## **References**

- Allison, I. et al. (2009). *The Copenhagen diagnosis: Updating the world on the latest climate science*. Sydney, Australia: The University of New South Wales Climate Change Research Centre (CCRC).
- Baker, M. (2016). Is there a reproducibility crisis? *Nature*, 533, 452-454.
- Biddle, J. (2016). Inductive risk, epistemic risk, and overdiagnosis of disease. *Perspectives on Science*, 24(2), 192-205.

- Biddle, J. & Kukla, R. (2017). The geography of epistemic risk. In: K. C. Elliott & T. Richards (Eds.), *Exploring inductive risk: Case studies of values in science* (pp. 215-237). Oxford University Press.
- Biddle, J. & Leuschner, A. (2015). Climate skepticism and the manufacture of doubt: Can dissent in science be epistemically detrimental? *European Journal for the Philosophy of Science*, 5, 261-278.
- Bird, A. (in press): Understanding the replication crisis as a base rate fallacy. *The British Journal for the Philosophy of Science*, ahead of print, <https://doi.org/10.1093/bjps/axy051>.
- Borgerson, K. (2009). Valuing evidence: Bias and the evidence hierarchy of evidence-based medicine. *Perspectives in Biology and Medicine*, 52(2), 218-233.
- Brown, M. J. (2013). Values in science beyond underdetermination and inductive risk. *Philosophy of Science*, 80, 829-839.
- Brysse, K., Oreskes, N., O'Reilly, J., & Oppenheimer, M. (2013). Climate change prediction: Erring on the side of least drama? *Global Environmental Change*, 23, 327-337.
- Carrier, M. (2010). Scientific knowledge and scientific expertise: Epistemic and social conditions of their trustworthiness, *Analyse & Kritik*, 32, 195-212
- Carrier, M. (2013). Values and objectivity in science: Value-ladenness, pluralism and the epistemic attitude. *Science and Education*, 22, 2547-68.
- Coady, C.A.J. (1992). *Testimony: A philosophical study*. Clarendon.
- Committee on Strategic Advice on the U.S. Climate Change Science Program; National Research Council (2009). *Restructuring federal climate research to meet the challenges of climate change*, Washington, D.C.: The National Academies Press.
- Daston, L. & Galison, P. (2007). *Objectivity*. Zone Books.
- Douglas, H. (2000). Inductive risk and values in science. *Philosophy of Science*, 67, 559-579.
- Douglas, H. (2004). The irreducible complexity of objectivity. *Synthese*, 138, 453-473.
- Douglas, H. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.
- Feest, U. (2019). Why replication is overrated. *Philosophy of Science*, 86, 895 -905.
- Fine, A. (1998). The viewpoint of no-one in particular. *Proceedings and Addresses of the American Philosophical Association*, 72, 9-20.
- Freudenberg, W. R. & Muselli, V. (2010). Global warming estimates, media expectations, and the asymmetry of scientific challenge. *Global Environmental Change*, 20, 483-491.

- Goldenberg, M. J. (2006). On evidence and evidence-based medicine: Lessons from the philosophy of science. *Social Science and Medicine*, 62(11), 2621-2632.
- Habermas, J. (1968/1971). *Knowledge and human interests*, trans. by Jeremy J. Shapiro. Beacon Press.
- Hardwig, J. (1991). The role of trust in knowledge. *The Journal of Philosophy*, 88, 693-708.
- Hartgerink, C. H. J., Wicherts, J. M. & van Assen, M. A. L. M. (2017). Too good to be false: Nonsignificant results revisited. *Collabra: Psychology*, 3(1), 9.
- Horkheimer, M. (1937/1999): Traditional and critical theory. In M. O'Connell (Ed.), *Critical theory: Selected essays* (pp. 188-243). Continuum Press.
- Intemann, K. (2015): Distinguishing between legitimate and illegitimate values in climate modeling. *European Journal for Philosophy of Science*, 5, 217-232.
- Ioannidis, J. P. A., Munafò, M. R., Fusar-Poli, P., Nosek, B. A., & David, S. P. (2014). Publication and other reporting biases in cognitive sciences: Detection, prevalence, and prevention. *Trends in Cognitive Sciences*, 18, 235-241.
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23, 524-532.
- John, S. (2015). Inductive risk and the contexts of communication, *Synthese*, 192, 79-96.
- John, S. (2019). Science, truth and dictatorship: Wishful thinking or wishful speaking? *Studies in History and Philosophy of Science*, 78, 64-72.
- Kourany, J. A. (2016): Should some knowledge be forbidden? The case of cognitive differences research. *Philosophy of Science*, 83, 779-790.
- Kuhn, T. S. (1977): Objectivity, value judgment, and theory choice. In: *The essential tension: Selected studies in scientific tradition and change* (pp. 320–339), University of Chicago Press.
- Latour, B. (2000): On the partial existence of existing and non-existing objects. In: L. Daston (Ed.), *Biographies of scientific objects* (pp. 247-269). University of Chicago Press.
- Lewandowsky, S., Oreskes, N., Risbey, J.S., Newell, B.R., & Smithson, M. (2015). Seepage: Climate change denial and its effect on the scientific community. *Global Environmental Change*, 33, 1-13.
- Longino, H. E. (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton University Press.
- Longino, H. E. (2002). *The fate of knowledge*. Princeton University Press.

- Megill, A. (1994). Introduction: Four senses of objectivity. In: A. Megill (Ed.), *Rethinking objectivity* (pp. 1-20). Duke University Press.
- Nozick, R. (2001). *Invariances: The structure of the objective world*. Belknap.
- Rahmstorf, S. (2007): A semi-empirical approach to projecting future sea-level rise. *Science*, 315, 368-370.
- Romero, F. (2019). Philosophy of science and the replicability crisis. *Philosophy Compass*, 14(11), e12633.
- Rudner, R. (1953). The scientist qua scientist makes value judgments. *Philosophy of Science*, 20, 1-6.
- Scheufele, D. A., Corley, E.A., Dunwoody, S., Shih, T.-J., Hillback, E. & Guston, D. H. (2007). Scientists worry about some risks more than the public, *Nature Nanotechnology*, 2, 732-734.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359-1366.
- Sosa, E. (2003). The place of truth in epistemology. In: M. DePaul & L. Zagzebski (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology* (pp. 155-179). Clarendon.
- Van Fraassen, B. (1980). *The scientific image*. Oxford University Press.
- Wilholt, T. (2009): Bias and values in scientific research. *Studies in History and Philosophy of Science*, 40, 92-101.
- Wilholt, T. (2013). Epistemic trust in science. *The British Journal for the Philosophy of Science*, 64, 233-253.
- Wilholt, T. (2016). Collaborative research, scientific communities, and the social diffusion of trustworthiness. In: M. S. Brady & M. Fricker (Eds.), *The epistemic life of groups: Essays in the epistemology of collectives* (pp. 218-233). Oxford University Press.
- Worrall, J. (2002). What evidence in evidence-based medicine? *Philosophy of Science*, 69(S3), S316-S330.