METAPHILOSOPHY


METAPHILOSOPHY
Vol. 50, Nos. 1–2, January 2019
0026-1068

# NORMATIVITY IN THE PHILOSOPHY OF SCIENCE

MARIE I. KAISER

**Abstract**: This paper analyzes what it means for philosophy of science to be normative. It argues that normativity is a multifaceted phenomenon rather than a general feature that a philosophical theory either has or lacks. It analyzes the normativity of philosophy of science by articulating three ways in which a philosophical theory can be normative. Methodological normativity arises from normative assumptions that philosophers make when they select, interpret, evaluate, and mutually adjust relevant empirical information, on which they base their philosophical theories. Object normativity emerges from the fact that the object of philosophical theorizing can itself be normative, such as when philosophers discuss epistemic norms in science. Metanormativity arises from the kind of claims that a philosophical theory contains, such as normative claims about science as it should be. Distinguishing these three kinds of normativity gives rise to a nuanced and illuminating view of how philosophy of science can be normative.

Keywords: normative, scientific practice, practice turn, philosophical methods, empirical information, epistemic norms.

## 1. Introduction

In recent decades, many areas of philosophy of science have undergone what is now referred to as a “practice turn” (e.g., Soler et al. 2014), that is, a turn towards scientific practice. More and more philosophers of science agree that philosophical theories about science must account for how science actually is done and must be informed, for instance, by the explanations developed in scientific practice and by the investigative strategies that scientists in fact employ. In other words, they agree that philosophical accounts about science must arise from an “empirical engagement with science” (Boumans and Leonelli 2013, 260) and that philosophers of science should seek to understand—“from the inside while retaining a philosophical perspective” (Wimsatt 2007, 27)—how science works, why it is successful, and why it sometimes fails. Philosophy of science that pays close attention to scientific practice is also called "philosophy of science in practice" (Ankeny et al. 2011; Mansnerus and Wagenknecht 2015).[1]

Philosophers spell out in different ways what it means for philosophy of science to turn towards scientific practice. Frequently, the practice turn is associated with a shift from analyzing products of science to analyzing scientific processes or activities (Chang 2011), with a shift from theory-focused to practice-centered epistemology (Waters 2014), and with a shift from purely intellectual and conceptual perspectives on science to ones that also consider the social context and the material aspects of scientific practice (Soler et al. 2014). Moreover, the turn towards scientific practice is simultaneously understood as a turn away from traditional normative theories about science, which construct ideals of how science should or ideally would work. These normative ideals were typically formulated ex cathedra and are thus criticized for being disconnected from and peripheral to the empirical reality of scientific practice. Accordingly, Soler and colleagues characterize the practice turn as a "shift from normative to descriptive perspectives on science" (2014, 15).

On the other hand, normativity continues to play a role in the philosophy of science after the practice turn. Several philosophers who pay close attention to scientific practice emphasize that science is an inherently collective activity and that we thus must take into account the social norms that influence the pursuit of scientific knowledge (Lloyd 2006; Kitcher 2011). In addition, some philosophers subscribe to the practice turn but sustain their normative aspirations. These philosophers use the results of their descriptive analyses to offer normative advice about how science should be done and about how certain concepts should be understood (e.g., Woodward 2003, 7; Craver 2007, viii). Finally, even if philosophers of science seek to describe a certain element of scientific practice, they seem to implicitly rely on normative assumptions, including assumptions about what are good examples and about how to assess the success of science. Hence, philosophy of science in practice seems to be thoroughly normative.

How can it be that the practice turn involves both moving away from normativity and at the same time leaving room for and even moving towards normativity? My goal in this paper is to solve this apparent contradiction. I argue that normativity in the philosophy of science is not a single matter, not a general feature that a philosophical account either has or doesn't have, but is a multifaceted phenomenon. The turn towards scientific practice involves moving away from one kind of normativity but not from others. I analyze in the paper the normativity of philosophy of science by articulating three ways in which a philosophical account can be normative. I distinguish metanormativity, methodological normativity,

[1] A recent indication of the turn towards scientific practice is the Society for Philosophy of Science in Practice (SPSP), founded in 2005.

and object normativity, and show how the different kinds of normativity relate to each other.

The structure of the paper is as follows. In section 2, I introduce the terminological framework of my analysis. Section 3 explicates the first way in which philosophy of science can be normative, which I call "metanormativity." I show that this type of normativity emerges from the kind of claims that a philosophical theory about some feature or element of science contains. I argue that philosophers of science who join the practice turn move away from metanormativity in the ex cathedra style, but that they need not abandon metanormativity in general. In Section 4, I examine the methodology of practice-oriented philosophy of science and reveal a second kind of normativity. My central claim is that methodological normativity arises from the need to make normative assumptions in selecting, interpreting, and evaluating the empirical basis of a philosophical account. In Section 5, I distinguish a third kind of normativity and argue that it arises from the fact that some philosophers of science discuss the role of epistemic and social norms in science. I refer to this kind of normativity as "object normativity" because it is due to the object of philosophical theorizing itself being normative. Finally, in section 6, I point out how the three kinds of normativity can be combined and how they depend on each other.

## 2. Terminological Framework

This section introduces the basic concepts on which my metaphilosophical analysis relies as well as the figure, to be successively refined in subsequent sections, that I use to illustrate my claims. The philosophy of science is a philosophical discipline that can be said to consist of different philosophical accounts or theories. You might prefer to think of philosophy as being made up of philosophical positions, philosophical questions and answers, or philosophical problems and solutions. My focus on philosophical accounts and theories does not exclude this. For the purposes of this paper, I use "philosophical account" and "philosophical theory" interchangeably and in a broad sense—though I am aware of the fact that some philosophers of science use "philosophical theory" in a stricter sense, for instance, as referring to sets of claims that are or consist of definitions that specify necessary and sufficient conditions.

A philosophical account or theory belongs to the philosophy of science if it makes claims about science. In other words, the object of philosophical theorizing in the philosophy of science is either some feature of science, such as systematicity, or a certain element of science, such as computer simulations, theoretical terms, model organisms, interventions, or causal inferences. We can say that a philosophical theory *T* concerns or is about some feature or element of science *E*. Figure 1 illustrates this way

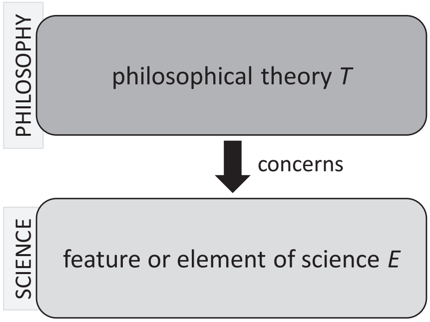


FIGURE 1. Philosophical theories in the philosophy of science

of conceptualizing the relation between philosophy of science and science.[2]

The expression "feature or element of science" is supposed to capture any entity that philosophers of science reason about. "Element of science" refers not only to material objects (for example, microscopes, electrons, scientists) but also to, for instance, scientific activities (such as modeling) and linguistic entities (such as observation statements). I speak of features and elements *of science*, not of scientific practice, because philosophers of science also make claims about hypothetical science, while "scientific practice" has the connotation of referring to actual science only (more on the notion of scientific practice in section 4). One might object that philosophers of science do not always make claims about science. Sometimes they are interested in metaphysical issues, which they take to be claims about the natural world itself rather than claims about how scientists investigate the natural world. I agree that we should not exclude metaphysical claims from the philosophy of science—especially if metaphysical questions are to be addressed in a naturalistic fashion (that is, by analyzing scientific knowledge). For reasons of simplicity, however, I will stick to the phrase "feature or element of science" and not always mention that philosophical theories might also concern the natural world studied by science.

## 3. Metanormativity

This section explicates the first of three ways in which philosophy of science can be normative. I argue that this kind of normativity, which I call "metanormativity," arises from the kind of claims that a philosophical

[2] Figure 1 is idealized in several respects. What is most important, it suggests that there exists a sharp boundary between philosophy (of science) and science, which is very difficult, if not impossible, to draw in reality.

theory contains. It is this kind of normativity that explains the manner in which the practice turn in the philosophy of science can be associated with a turn away from normative perspectives on science.

Let us start with an episode in the history of philosophy of science. The debate about epistemic reduction in the biological sciences started with the attempt to apply Nagel's (1961) formal model of theory reduction to biological cases, such as the putative reduction of Mendelian genetics to molecular biology (e.g., Hull 1974; Rosenberg 1985). It quickly became clear that Nagel's model encounters serious obstacles when applied to the biological sciences. As a response, Schaffner developed Nagel's model further and proposed his General Reduction-Replacement (GRR) model (1974, 1993). Schaffner explicitly constructs the GRR model as an ideal that need not be realized in contemporary scientific practice to be correct. He admits that the GRR model is only "peripheral" (1974, 111; 1993, 509) to biological practice because molecular biologists are not interested in obtaining the "complete chemical characterizations" (1974, 127) that are required for the kind of theory reductions he envisions. He treats the GRR model as a regulative ideal that should, but does not in fact, guide the development of molecular biology (1993, 511).[3] Other philosophers of science took the obstacles to applying Nagel's model to biology as evidence for antireductionism (Waters 1990). Nevertheless, almost all philosophers of biology agreed that Nagel's model was an adequate view of epistemic reduction in biology; at the time, it sounded "suspicious to change the standards of reduction" (Rosenberg 1985, 110). The situation changed in the 1990s when more and more philosophers realized that it does not make sense to impose an ill-fitting ideal of reduction on the biological sciences. Since then, several philosophers of biology have developed alternative accounts of epistemic reduction, which are based on extensive analyses of cases of epistemic reduction that actually occur in biological practice (for example, reductionist heuristics [Wimsatt 2006, 2007; Waters 2008], and reductive explanations [Sarkar 1998; Hüttemann and Love 2011; Kaiser 2015]). The debate about epistemic reduction in biology is only one example where the turn towards the empirical reality of scientific practice was accompanied by a turn away from philosophical accounts that construct ex cathedra normative claims about how science should be pursued or what it ideally looks like.

[3] In his recent work, Schaffner concedes that "what have traditionally been seen as robust reductions of one theory or one branch of science by another more fundamental one are largely a myth" (2006, 378). At first sight, this seems as an immense departure from his original position. Under closer inspection, however, one notices that Schaffner still regards the GRR model as an "ideal" (2006, 384) of what a complete reduction in biology would look like. For instance, he argues that in biology reductive, causal mechanical explanations are mere "*partial* reductions" and "reductions of the *creeping* sort" (2006, 397; emphasis in the original). This argument presupposes that there is an ideal of a complete, fully satisfying reduction.

### *3.1. Normative Claims About Science*

What makes philosophical accounts, such as Schaffner's GRR model, normative? The object of Schaffner's account is epistemic reduction in biology (and medicine). This is the element of science *E* that the GRR model is concerned with. Schaffner does not describe what epistemic reduction in biology in fact is. Rather, he expresses how philosophers should understand the concept of epistemic reduction, states what good cases of epistemic reduction in biology are, and argues that biologists should try to achieve epistemic reductions that satisfy the requirements specified in the GRR model. His model is normative because it contains claims about its object of study (that is, epistemic reduction in biology), which are normative (rather than factual). We can generalize this thesis so that it holds for all philosophical accounts in the philosophy of science.

> *Metanormativity.* A philosophical theory *T* about a feature or element of science *E* is metanormative iff *T* contains normative claims about *E*.

This is the first way in which philosophy of science can be normative. I call this kind of normativity "metanormativity" because it arises from a feature of the philosophical account itself, rather than from how the account is developed or from what it is about. Figure 2 illustrates this kind of normativity.

What does it mean for a philosophical theory to contain normative claims about a feature or element of science? In general, normative claims can be evaluative statements and express the fact that something has or lacks a certain value, that something is good or bad, correct or incorrect. Normative claims can also be prescriptive and offer advice about what ought or ought not to be the case (for the distinction between evaluative and prescriptive norms see, e.g., McHugh 2012). In the philosophy of

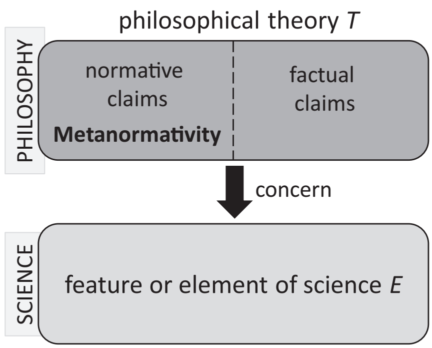


FIGURE 2. Metanormativity

science, normative claims about a certain feature or element of science are often normative in both the evaluative and the prescriptive sense. Schaffner's GRR model, for example, is metanormative in both senses. On the one hand, Schaffner's model contains the evaluative normative claim that only epistemic reductions that satisfy the GRR model are good or correct, other epistemic reductions being of the "*creeping* sort" (2006, 397; emphasis in the original) and thus lacking value. On the other, the GRR model is prescriptively metanormative because Schaffner assumes that philosophers should conceive of epistemic reduction according to his GRR model and that the model should guide biological research. Other examples of metanormative philosophical theories include Popper's (1959) view that falsifiability is the benchmark of science, Hempel's theory of what scientific explanation should look like (which is "not meant to describe how working scientists actually formulate their explanatory accounts" [1965, 412]), and Brandon's "normative ideal" (1996, 197) for adaptation explanations in evolutionary biology.

Normative claims are usually contrasted with factual (or positive or descriptive) claims that attempt to describe reality and thus are truth apt. Examples of philosophical theories that consist of only factual claims about their objects of study are Waters's (2008) analysis of investigative strategies in genetics and Winther's (2011) account of the integration of different kinds of part-whole explanations. These philosophical accounts are not metanormative, because they contain only factual claims about the element of science that they analyze. Waters describes how the central investigative strategy of classical genetics, the genetic approach, in fact works. Winther explicates how scientists actually develop and integrate different kinds of part-whole explanations of the tetrapod limb. Both describe science as it actually is, instead of making claims about science as it should be or about what is good science.

To conclude, one way in which philosophy of science can be normative is that it can make normative claims (that is, express evaluations and offer advice) about the feature or element of science that it studies. I call this kind of normativity "metanormativity" because it arises from a feature of the philosophical theory itself, namely, from the kind of claims that the philosophical theory contains.

### *3.2. Metanormativity After the Practice Turn*

Having clarified the concept of metanormativity, I can now examine whether the practice turn in the philosophy of science is a turn away from metanormativity. If philosophical theories about science take into account the empirical reality of scientific practice, does that imply that they are not and cannot be metanormative? In what follows, I argue that the practice turn involves only a turn away from a specific style of metanormativity, which I call "ex cathedra metanormativity," but that philosophy of science after the practice turn is compatible with other forms of metanormativity.

Philosophers of science who call for abandoning normative perspectives on science (e.g., Soler et al. 2014, 15) typically direct their criticism towards a very specific kind of metanormativity. They criticize philosophers who adopt a privileged viewpoint outside science and tell the practicing scientists what good science really is and how science is properly done. The objection is that these philosophers act as if they had "epistemic sovereignty" (Rouse 2002, 180), were "philosopher kings or philosophical police" (Sober 2008, xv), and were allowed to make metanormative claims about what is good science and how science should work ex cathedra (or from the ivory tower) without paying attention to how science in fact works. It should be emphasized that this is a caricature which very few, if any, philosophers of science fulfill. In a less radical version, metanormativity is ex cathedra if metanormative claims about what is good science and about how science ought to be pursued are developed and justified without taking into account actual scientific practice (this is what McMullin characterizes as "external philosophy of science" [1970, 24]). In other words, ex cathedra metanormative claims about a certain feature or element of science are not informed by and cannot fail in light of the empirical reality of scientific practice.

Examples of ex cathedra metanormativity can be found, for instance, in the metaphysics of science literature. Best-systems accounts of laws of nature (e.g., Lewis 1999, chaps. 1 and 15) state that generalizations in science should be regarded as laws only if they appear as axioms or theorems in the best (that is, the simplest and strongest) deductive system that contains everything we know in terms of natural properties. The metanormative claims in best-systems accounts are ex cathedra because they are said to be adequate independently of whether they capture actual cases of laws from scientific practice, whether they make sense of how scientists use the term "law," and whether all scientific knowledge can in fact be appropriately organized in the form of deductive systems and natural properties. That is, best-systems accounts of laws of nature exhibit ex cathedra metanormativity because they are not informed by and cannot fail in light of the empirical reality of scientific practice.

Ex cathedra metanormativity is thus not compatible with philosophers turning their attention to scientific practice, because "ex cathedra" means exactly the opposite, namely, ignoring scientific practice in developing and justifying a philosophical account. In contrast, metanormativity that is *not* in the ex cathedra style is perfectly compatible with the practice turn. Some philosophers even claim that metanormativity is an indispensable feature of *any* philosophical theory about science (e.g., Wimsatt 2007, 26; Sober 2008, xv).

What does it mean for philosophy of science to evaluate science and offer advice that is "contextual and sensitive to feedback" (Wimsatt 2007,

27) and that is informed by actual scientific practice? In my terminological framework, this means that a philosophical theory about a certain feature or element of science *E* contains only such normative claims about *E* that take into account, are drawn from, or are informed by factual claims about *E*. For example, Craver states that his descriptive project of characterizing mechanistic explanations in neurosciences "is the first step in a normative project: to clarify the distinction between good explanations and bad" (2007, viii). Similarly, Woodward emphasizes that his theory of causality and causal explanation also "makes recommendations about what one ought to mean by various causal and explanatory claims, rather than just attempting to describe how we use those claims" (2003, 7).

An interesting issue that I can only touch on here concerns the relation between normative and factual claims. Philosophers who aim at offering advice and evaluating science while paying close attention to the empirical reality of scientific practice face a challenge. On the one hand, they can only avoid an ex cathedra stance if they keep the relation between normative and factual claims about science as close as possible. Metanormative claims about some feature or element of science *E* should not be developed and justified independently from factual claims about *E*. Instead, metanormative claims about *E* should be based on or informed by factual claims about *E*. Simply deriving normative claims from factual claims, however, is illegitimate because it amounts to an is-ought fallacy (Bechtel and Richardson 2010, 10). On the other hand, philosophers of science might want to avoid the is-ought fallacy by developing and justifying their metanormative claims completely independently from factual claims about scientific practice (for example, by adducing a priori reasons [Schindler 2013]). If philosophers of science do this, however, their advice and evaluations become detached from the empirical reality of scientific practice, and the bugaboo of ex cathedra metanormativity looms again. To conclude, philosophers of science who have undergone the practice turn and still make metanormative claims need to meet this challenge and find ways to link their metanormative claims closely—but not too closely—to their factual claims. Promising approaches make use of, for example, the idea of a reflective equilibrium to specify how normative conclusions can be drawn from descriptive matters (Thagard 1988, chap. 7; cf. van Thiel and van Delden 2010).

## 4. Methodological Normativity

In this section, I analyze the methodology of philosophy of science in practice (PSP). My central claim is that even a philosophy of science that seeks to understand and accurately describe a certain element of scientific practice, and that thus contains only factual claims about science, is thoroughly normative. I call this kind of normativity "methodological normativity" because it arises from the fact that the methodology of PSP inevitably involves normative assumptions.

### *4.1. The Methodology of Philosophy of Science in Practice*

So far in this paper, we have been given only a rough idea of what it means for philosophy of science to pay close attention to actual scientific practice. Proponents of PSP have in common that they seek to understand science "from the inside while retaining a philosophical perspective" (Wimsatt 2007, 27), and that their philosophical accounts arise from an "empirical engagement with science" (Boumans and Leonelli 2013, 260). What do statements like these imply for the methodology of PSP?

Philosophy of science is a second-order discipline that studies the sciences that, in turn, study the natural world (e.g., McMullin 1970, 27; Sober 2008, xv).[4] For example, medical scientists aim at discovering the causes of complex diseases, such as cancer. Philosophers of science, in contrast, seek to understand, for instance, causal reasoning in cancer science as well as the strategies that cancer scientists employ to deal with causal complexity. In order to understand the methodology of PSP it is helpful to see in how far it presupposes a minimal methodological naturalism (cf. Giere 1999, 53–54; Bechtel 2008, 4–10). In my view, the methodology of PSP is similar to scientific methodology in at least one minor respect: in both fields, the theory or account that is developed must be empirically adequate, that is, it must capture and find evidential support in the available empirical data. This is why PSP is characterized as "empirical philosophy of science" (Mansnerus and Wagenknecht 2015, 38). Among the differences between the natural sciences and PSP is that empirical data in the former are about the natural world, whereas empirical data or information in the latter is about the natural sciences. That is, the empirical information against which a philosophical theory is "tested" is information from and about scientific practice. Scientists, in turn, develop scientific theories that they test against empirical data about the natural world (cf. Paul 2012).[5]

The claim that philosophy of science after the practice turn involves an empirical engagement with science can thus be specified as follows. In PSP, empirical information from and about scientific practice plays a central role in developing and justifying any philosophical account or theory. For example, when developing a philosophical account of causal inference

[4] This holds even for scientific metaphysics (e.g., Ross, Ladyman, and Kinkaid 2013), which studies scientific knowledge (for example, scientific theories or successful scientific practices) to draw metaphysical inferences.

[5] It is compatible with minimal methodological naturalism that the methodology of PSP and the methodology of the natural sciences differ in other respects. For example, one might claim that PSP is not a "science of science," because it is hermeneutic and proceeds through "acts of interpretation" (Schickore 2011, 461).

in oncology philosophers must take into account empirical information about the interventionist studies that scientists perform to identify carcinogens, about the causal explanations that oncologists develop, about the strategies that they employ to distinguish mere correlations from causal relations, and the new insights into the progression of cancer that are gained in these studies.

Empirical information that can be relevant to philosophical analysis is of diverse kinds. Accordingly, I understand the notion of scientific practice in a broad sense. It encompasses elements that are material or practical in a stricter sense (for example, scientific instruments, experiments, lab conditions, methods, and model organisms) as well as elements that are more theoretical, such as a scientist's epistemic activities (for example, explaining, testing, observing, modeling, theorizing, idealizing) and the results of these activities (for example, explanations, models, theories, generalizations).[6] Empirical information from and about scientific practice can also be information about the history of scientific practices (for example, in the form of historical case studies), which can lead to an integrated history and philosophy of science (cf. McMullin 1970; Schickore 2011; Kinzel 2015).

What does it mean for empirical information to play a central role in developing and justifying a philosophical theory? Philosophical accounts that pay close attention to scientific practice consist of factual claims that describe how science in fact works (in addition, they may contain metanormative claims as well; recall section 3). Unlike what the word "describes" suggests, philosophical accounts cannot be pure descriptions or one-to-one mappings of scientific practice. A philosophical account that is coherent and provides clarity and understanding cannot simply be read off scientific practice. Rather, it must result from a critical reconstruction of relevant empirical information from scientific practice. This is what Wimsatt seems to have in mind when he emphasizes that philosophical accounts must be developed from the inside of science while retaining a philosophical perspective (2007, 27), and this is what distinguishes PSP, for instance, from science journalism. In my view, the process of critically reconstructing relevant empirical information involves four major tasks: first, *selecting* empirical information from and about scientific practice that is relevant; second, *interpreting* empirical information, for example, by abstracting from irrelevant details and making explicit underlying assumptions; third, critically *evaluating* empirical information with the aim of establishing coherence; and fourth, *mutually adjusting* philosophical claims and empirical information until a reflective equilibrium is reached.

[6] Some authors put forward a narrower notion of scientific practice that includes material aspects of science only (Soler et al. 2014, 18) or that focuses on investigative practices (Waters 2014). In my view, theories and concepts remain important elements of scientific practice, which is why we should not exclude them.

I examine these four tasks in turn before addressing the question of which kind of normativity is involved in the process of critically reconstructing relevant empirical information from and about scientific practice.

First, the amount of empirical information that is available about most elements of scientific practice is enormous. Because of limited resources philosophers are forced to restrict their analyses to the empirical information that they think is particularly relevant to the question they address. The reasons why philosophers might pick out empirical information as relevant vary, but there are some general principles that guide the selection process. Typically, philosophers regard examples as relevant because they are paradigmatic or because they are representative for other cases of the same kind. By focusing on representative cases, philosophers try to make sure that the philosophical theory they develop holds not only for the analyzed cases but also for science or a scientific field in general. Another major reason for assessing a case as relevant is that it is of particular importance to a scientific field, for instance, if it has driven scientific research for a longer period of time or if it is an example for how the success (or failure) of research in that field is promoted. Since success is seen as a central goal of science, it is also of particular interest to philosophers (Giere 1999, 53; Norton 2003, 648). Successful examples might, for instance, be those that appear in established textbooks or that are much discussed in a certain field. A case might also be assessed as relevant because it contributes to achieving another central goal of a scientific field, such as manipulation or disease control, or because it concerns central processes of life or of our world (for example, reproduction of living beings, quantum entanglement, and the Big Bang).

Second, the process of developing a philosophical theory while taking relevant empirical information into account is often not straightforward but involves a great deal of abstraction, explication, and "interpretation" (Schickore 2011, 471). Philosophers must explicate background assumptions that scientists implicitly presuppose in their experimentation and reasoning, they must establish connections between seemingly unrelated claims and concepts, they must abstract from philosophically irrelevant details, and they must draw philosophical inferences from empirical information. Consider the example of developing an account of what makes biological explanations reductive (Hüttemann and Love 2011; Kaiser 2015). The first challenge that philosophers encounter is that only very few biologists indicate whether the explanations they give are reductive or not. They argue about the adequacy of explanations but not about their reductive character because this is just not important to them. Some biologists engage in intensive debates about the "limits of reductionism" (Mazzocchi 2008, 10) and the need to move "beyond reductionism" (Gallagher and Appenzeller 1999, 79). But even then biologists rarely speak about reductive explanations. Rather, they discuss the correctness of a reductionist approach and the adequacy of applying reductive methods. Sometimes biologists do not even use terms containing “reduc-” but, nevertheless, express assumptions about reductive explanation, such as when they discuss part-whole explanations or the method of decomposition. This example illustrates the extent to which the process of developing philosophical claims on the basis of empirical information requires acts of abstraction, explication, and interpretation.

Third, philosophers are confronted with many differences or even inconsistencies within and among scientific fields. For instance, explanatory and investigative strategies vary, the same concepts are understood differently, and different background assumptions are made. To develop a coherent, unified theory about a certain feature or element of scientific practice, philosophers must take up a critical stance and sort out the empirical information that can be dismissed as false, misleading, or biased. For example, what biologists mean by “reductionism” and what they think constrains the adequacy of a reductive explanation is by no means homogenous and involves inconsistencies. Some biologists identify reductive explanation with additive explanations, that is, with explanations in which biological systems are treated as aggregative systems (Kitano 2002, 1662). Other biologists explicitly reject this claim because it results in a too restricted view of reductive explanation. They state that “[m]olecular biologists . . . do not hold the naive view that complex structures and processes are just sums of their parts” (Fincham 2000, 343). If one wants to develop a coherent theory of reductive explanation, one needs to ponder which of these claims should inform the philosophical theory (for example, because they are more common or are best in line with other relevant empirical information) and which should be sorted out as incorrect, rare, too vague, or insufficiently justified.

Finally, as the other three tasks already indicate, the process of developing a philosophical theory by taking into account empirical information from and about scientific practice is not a one-way process but involves a repeated mutual adjustment and moving back and forth between philosophical theory and empirical information (this is why figure 3 includes arrows leading from science via empirical information to philosophy and back again). This process can also be characterized as an inherently hermeneutic endeavor (Schickore 2011) and as an iterative (inner) dialogue between abstract theory and concrete data (Mansnerus and Wagenknecht 2015). The process of mutual adjustment often starts with provisional philosophical claims and preconceptions that are brought together with, sharpened, and modified in the light of provisional selections, interpretations, and evaluations of empirical information. The process comes to an end, for example, as soon as a reflective equilibrium between philosophical theory and empirical information is reached (Thagard 1988, 119; for the general idea see, e.g., Elgin 1996, chap. 4).

To conclude, in the philosophy of science after the practice turn, developing a philosophical theory *T* about some feature or element of science *E*

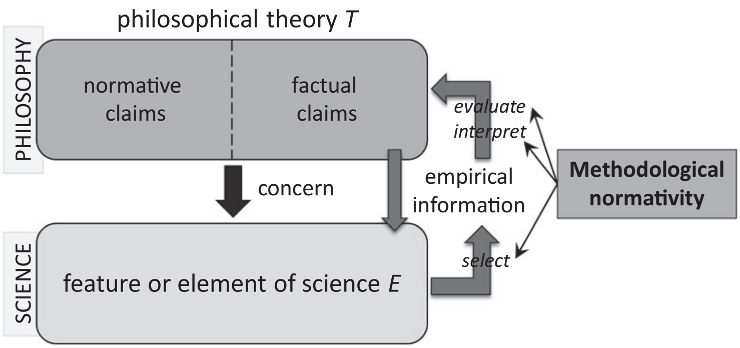


FIGURE 3. Methodological normativity

involves selecting, interpreting, and evaluating relevant empirical information about *E* from and about scientific practice as well as making mutual adjustments between philosophical theory and empirical information.

*4.2. Methodological Normativity*

The goal of this section is to show that even empirically based philosophy of science, which seeks to understand and describe how scientific practice in fact works, is thoroughly normative. The kind of normativity involved here is different from metanormativity because the normative claims do not concern the object of philosophical theorizing itself. Rather, they concern the methodology by which an empirically based philosophical theory is developed. This is why I refer to this kind of normativity as "methodological normativity." Figure 3 illustrates this kind of normativity.

If philosophers develop a philosophical theory by selecting, interpreting, and evaluating relevant empirical information and mutually adjusting philosophical claims and empirical information, they presuppose, usually implicitly, certain methodological norms. These are norms that, for instance, express what is good empirical information (relative to the philosophical question at stake) and that give advice about how philosophers should proceed in developing a particular theory through selecting, interpreting, and evaluating empirical information and mutually adjusting theory and empirical information. The following examples of methodologically normative assumptions illustrate what these methodological norms are that guide theory development in PSP:

An example of *E* is good because it is an instance of successful science/contributes to a major aim of a scientific field (for example, manipulation, disease control, prediction, or technological progress).

An example of $E$ is good because it is paradigmatic/clear-cut/robust/representative for cases of the same kind.
Empirical information about $E$ from a certain scientific field is good because this field is especially successful.
Empirical information about another element $E^*$ should be taken into account by a philosophical theory $T$ about $E$ because $E$ and $E^*$ are closely related.
Claims of scientists about $E$ should be excluded from the empirical basis of $T$ because these claims are incorrect/biased/too vague.
Empirical information about $E$ should not be taken into account by $T$ because it would prevent establishing coherence.
A claim of philosophical theory $T$ should be revised/abandoned because it conflicts with relevant empirical information about $E$.

As presented above, some of these methodological norms are formulated as evaluative normative claims, others as prescriptive. Nevertheless, all claims seem easily translatable from one formulation to another. This reflects the fact that methodological norms that figure in philosophy of science after the practice turn are typically both evaluative and prescriptive. That is, they express that certain kinds of empirical information or philosophical procedures have or lack a value, and they offer advice about what philosophers ought or ought not to do (for example, which empirical information they should take into account and which they should ignore).

In most cases, methodologically normative assumptions will not be explicitly stated but rather will implicitly guide the selection, interpretation, and evaluation of empirical information and the process of mutually adjusting philosophical theory and empirical information. Furthermore, we need not assume that these methodological norms are static. They can change over time, if we, for instance, learn how to better assess or select empirical information from scientific practice.

To conclude, philosophical theories that make factual claims about a certain feature or element of science may not be metanormative (if they contain factual claims only; recall section 3). Still, these theories are normative because the philosophical methodology of developing factual claims about science while taking into account empirical information from and about scientific practice inevitably involves making (implicit) normative assumptions about how to select, interpret, and evaluate empirical information and how to mutually adjust theory and empirical information. I call this second kind of normativity "methodological normativity" because it arises from the methodology by which an empirically based philosophical theory about science is developed.

*Methodological normativity.* A philosophical theory $T$ about a feature or element of science $E$ is methodologically normative iff $T$ contains factual claims about *E* that (implicitly) rely on normative claims about how to select, interpret, evaluate, and mutually adjust empirical information about *E* in developing *T*.

Methodological normativity differs from metanormativity in two important ways. First, methodological normativity stems from norms that concern the philosophical methodology and that thus commit philosophers, not scientists, to handling empirical information in a certain way or to seeking coherence between philosophical theory and empirical information in a certain way. Metanormativity, in contrast, commits primarily scientists to, for instance, seeking theory reductions à la Schaffner or adopting a specific understanding of the concept of a mechanism. Second, methodological normativity does not require that a philosophical theory contains normative claims about its object of study *E* (as metanormativity does). Rather, methodological normativity applies to empirically based philosophy of science only, that is, to philosophical theories that contain factual claims about *E*. If a philosophical theory contains factual and normative claims about *E* it is methodologically normative and metanormative (see section 3.2). In sum, introducing the category of methodological normativity reveals in what way even a philosophical theory that describes how science in fact works is thoroughly normative.

## 5. Object Normativity

In this section, I identify a third way in which philosophy of science can be normative: object normativity. This kind of normativity emerges from the fact that the object of philosophical theorizing itself can be normative. This is the case if philosophers reason about epistemic or social norms and their roles in science.[7] Among the questions that are of philosophical interest is, for instance, the question of whether epistemic norms such as simplicity, precision, explanatory power, and predictive success guide how scientists identify their objects of study, interpret empirical data, and choose between competing theories or explanations (e.g., Kuhn 1962). Philosophers of science also controversially discuss whether social or political norms, such as democracy, human rights, or gender biases, influence the scientific process of acquiring knowledge about the natural world and may jeopardize the objectivity of scientific knowledge (Longino 1990; Kitcher 2011). Because this kind of normativity arises from the objects of a philosophical theory being norms (or being related to norms) I refer to it as "object normativity."

[7] I speak about norms in science, rather than about values, because the concept of a norm is broader and accounts for evaluative as well as for prescriptive normative claims in science. I understand "social norms" in a broad way including various kinds of non-epistemic norms that are relevant to society.

*Object normativity.* A philosophical theory *T* about a feature or element of science *X* is object normative iff *T* refers to epistemic or social norms in science.

In the following two subsections, I explicate what it can mean for a philosophical theory to refer to epistemic or social norms in science and thereby introduce different subtypes of object normativity. First, I distinguish philosophical theories that describe which norms are in fact accepted in scientific practice from philosophical theories that posit norms that should be accepted in science (section 5.1). Second, among philosophical theories that refer to epistemic norms in science, I distinguish those that concern epistemic norms themselves from philosophical theories that are about some nonnormative element of science but that relate this element to certain epistemic norms (section 5.2).

Before moving on let me add a final general remark. One might wonder how object normativity relates to the turn towards philosophical accounts that engage with actual scientific practice. The practice turn is said to involve a shift to perspectives on science that are "more realistic" (Soler et al. 2014, 18), for instance, because they recognize the deep intertwinement of science and society. In line with this, Rouse warns against construing science as "clearly bounded and distinct from extrascientific 'context'" (2002, 164). Philosophers of science in practice typically avoid this danger because they take into account the practice of scientific research in its entire variety, including processes of inquiry, institutional settings, and social dynamics among investigators (Boumans and Leonelli 2013). This includes, for example, recognizing the political dimension of knowledge and the ways in which scientific fields may be shaped by the uses to which scientific knowledge may be put, such as how gene patents affect medical genetic testing (Carrier, Howard, and Kourany 2008).

Accordingly, some philosophers of science in practice argue that examining the goals of scientific activities requires not only epistemological considerations but also reflections on "the values, norms, and ideals inherent in the pursuit of scientific knowledge" (Ankeny et al. 2011, 305). This might suggest that any philosophical theory in the philosophy of science in practice *must* refer to epistemic or social norms and thus be object normative. I think that this claim is too strong and that we gain nothing from imposing such a strict requirement on what is seen as "proper" practice-oriented philosophy of science. Nevertheless, from a general perspective, the practice turn is accompanied by a shift to a philosophy of science that is more object normative.

### *5.1. Describing Norms Versus Positing Norms*

The first distinction of subtypes of object normativity results from linking object normativity to metanormativity. In Section 3, I distinguished philosophical theories that contain normative claims about their objects of study *E* (and thus are metanormative) from those that contain only factual claims about *E*. This difference applies to any feature or element of science—to nonnormative elements of science, such as causal inferences, model organisms, and mathematical equations, as well as to epistemic and social norms that figure in science (which give rise to object normativity), such as explanatory power or gender biases. Figure 4 illustrates the four possible types of philosophical theories that result from combining the two different kinds of claims that a philosophical theory can contain with the two kinds of objects that the theory can be concerned with.

Combinations ① and ② represent philosophical theories that are not object normative, because they refer to nonnormative elements of science only. Philosophical theories of type ① contain normative claims about nonnormative elements of science (for example, the claim that only cases of reductions that fulfill Schaffner's GRR model are good) and thus are metanormative. By contrast, philosophical theories of type ② make only factual claims about nonnormative elements of science and thus are neither metanormative nor object normative. The difference between ① and ② was spelled out in detail in section 3.

Consider now philosophical theories that refer to epistemic or social norms in science and thus are object normative (combinations ③ and ④). Combination ④ represents philosophical theories that describe which epistemic or social norms are in fact accepted in scientific practice. These philosophical theories make factual claims about which norms actually influence scientific inquiry. I refer to this as "describing norms." Examples of descriptions of norms in science are Lloyd's (2006) theory of how gender biases influence the development of adaptive explanations of female orgasm and my analysis of how biologists evaluate reductive explanations

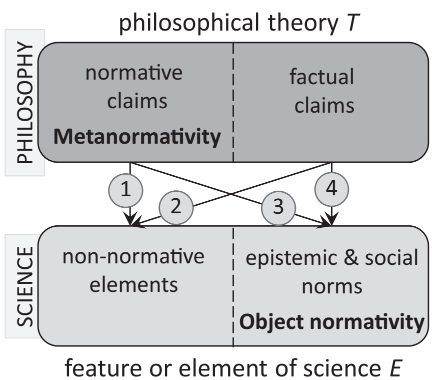


FIGURE 4. Possible combinations of metanormativity and object normativity

as adequate, that is, which norms of reductive explanation they accept (Kaiser 2015). By contrast, combination ③ represents philosophical theories that posit which epistemic or social norms should apply to science and justify why these norms should be accepted. I refer to this as "positing norms." Philosophical theories of this kind refer to norms by making normative claims about which epistemic or social norms ideally should influence, for instance, the process of gaining scientific knowledge (and how they should do so). Examples of normative claims about norms in science include: Kitcher's theory of well-ordered science, which constructs an idealistic picture of how decisions about the significance of research projects should be democratically assessed (2011); philosophical theories of simplicity as the best criterion for choosing among competing scientific theories and explanations (Sober 1975; Thagard 1988; White 2005); and Craver's (2007) account of how mechanistic explanations in neuroscience should be evaluated. In sum, philosophical theories that describe accepted norms (that is, ④) are as object normative as philosophical theories that posit which norms should be accepted (that is, ③), but only the latter are also metanormative because only they contain normative claims about the norms to which they refer.[8]

### *5.2. Theorizing About Norms Versus Relating to Norms*

In this section, I reveal a distinction between two subtypes of object normativity that applies only to philosophical theories that refer to *epistemic* norms. This distinction emerges from the fact that claims about epistemic norms in science—whether factual or normative—can figure differently in a philosophical theory. On the one hand, epistemic norms can be the objects of philosophical theorizing. In these cases, the philosophical theory is *about* these norms. For example, Thagard (1988) has proposed a theory about the epistemic value of simplicity, considering how we should understand it and why it is justified. Similarly, Lloyd's (2006) analysis is about gender biases and how they affect the development of adaptive explanations of female orgasm. I refer to these cases as "theorizing about norms."

On the other hand, reference to epistemic norms can be less central to a philosophical theory. In these cases, the object *E* of a philosophical theory is not epistemic norms but another nonnormative element of science (for example, causal inference, reduction, or the concept of a gene), and the philosophical theory includes claims about how *E* is related to certain epistemic norms (that either are in fact or should be accepted in science).

[8] One might express the difference between ③ and ④ also by claiming that philosophical theories of type ④ refer to intrinsic norms (that is, norms that are inherent in scientific practice), whereas those of type ③ refer to extrinsic norms (that is, norms that are posited from outside science).

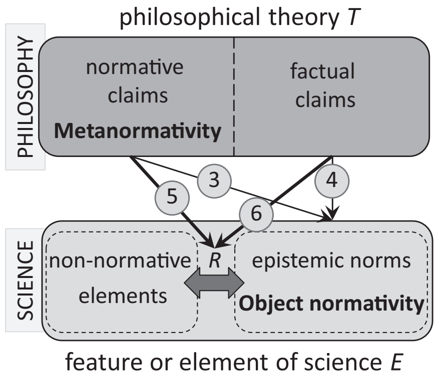


FIGURE 5. Subtypes of object normativity[9]

For example, philosophical analyses of the concept of reductive explanation are about reductive explanations but they may also elucidate why reductive explanations are adequate with respect to some phenomena and inadequate with respect to others, that is, how the reductive character of an explanation affects the epistemic value of explanatory success. This is an example of what I refer to as "relating to norms."[10] Figure 5 illustrates the difference between object normativity as theorizing about norms and as relating to norms.

Object-normative philosophical theories can either directly address epistemic norms in science, that is, theorize about norms (③ and ④), or they can address another element of science *E* and its relation *R* to certain epistemic norms, that is, they can relate *E* to epistemic norms (⑤ and ⑥). This is the difference between object normativity as theorizing about norms and as relating to norms. These two subtypes of object normativity can be combined with the other two subtypes of object normativity, which I introduced in section 5.1, namely, describing norms and positing norms. Object-normative philosophical theories can either make factual claims

[9] Philosophical theories that do not refer to epistemic or social norms (① and ② in figure 4) are not object normative and thus are omitted from figure 5.

[10] To which of these two categories a philosophical theory belongs depends on how finegrainedly it is individuated. For instance, Craver's account of mechanistic explanation is not about epistemic norms. Still, his account includes claims about what the norms of mechanistic explanations are and should be (Craver 2007, 20 and 111). If we conceptualized these claims as a separate theory, it would be a case of theorizing about norms, otherwise Carver's account of mechanistic explanation would be a case of relating to norms.

about epistemic norms or about the relations to epistemic norms (④ and ⑥), or they can make normative claims about these norms or relations (③ and ⑤), which gives rise to metanormativity (see section 5.1).

Recognizing the difference between object normativity as theorizing about norms and as relating to norms enables us to assess whether and to what extent object normativity is or should be prevalent in practice-oriented philosophy of science. Some philosophers hold the apparently strong view that any philosophical theory about science should refer to epistemic norms and thus should be object normative. For instance, Waters argues that philosophers of science should provide an understanding of "how the sciences work (and don't work) with respect to epistemic virtues that we value" (2004, 48). A philosophical theory of what genes are, for example, should capture not only how scientists reason about genes, investigate genes, and use gene terminology. It also must clarify how the gene concept is related, for instance, to the epistemic norm of explanatory power (for example, by discussing the limitations of gene-based explanations) or to the epistemic norm of investigative utility (for example, by revealing the usefulness of chief methods in gene-centered sciences). At first sight, the claim that any philosophical theory about science must refer to epistemic norms seems to result in a too restrictive view of what proper philosophy of science is. If we apply the distinction between relating to norms versus theorizing about norms, however, we see that the claim is weaker and more plausible. Waters does not call for a philosophy of science that studies epistemic norms in science only (which would be object normative as theorizing about norms). Besides epistemic norms, there are plenty of other features and elements of science that are worthy of philosophical investigation. Waters's claim is that any philosophical theory about these other features and elements of science (such as gene-based explanations) must explicate how they relate to epistemic norms. Regardless of whether one thinks that Waters's claim that object normativity should be prevalent in the philosophy of science in practice is fully convincing or not, my analysis shows that this claim must be understood to concern a specific subtype of object normativity, namely, object normativity as relating to norms.

## 6. Interrelations Between the Three Kinds of Normativity

The goal of this section is to explicate how the three kinds of normativity that characterize the philosophy of science relate to each other and can be combined.

Depending on the kinds of claims that compose a philosophical theory and depending on the kinds of objects with which a philosophical theory is concerned, a philosophical theory falls into one of four groups. It can be (1) metanormative (if it makes normative claims about non-normative elements of science), (2) object normative (if it makes factual claims about epistemic or social norms in science), (3) metanormative and object normative (if it makes normative claims about epistemic or social norms in science), or (4) neither metanormative nor object normative (if it makes only factual claims about nonnormative elements in science) (see section 5.1). All four combinations are possible because metanormativity and object normativity are not only distinct types of normativity but also independent of each other. The kind of claims that a philosophical theory makes does not commit it to a specific kind of object of philosophical theorizing (and vice versa).

Now, whenever a philosophical theory contains factual claims about some feature or element of science, it is methodologically normative because developing these factual claims on the basis of empirical information from and about scientific practice presupposes (implicit) normative claims about how to select, interpret, evaluate, and mutually adjust empirical information (see section 4.2). This holds for factual claims about nonnormative elements of science as well as for factual claims about norms in science. Hence, methodological normativity can but need not be combined with object normativity. It also holds for philosophical theories that contain not only factual claims but also normative claims (see section 3.2). The only kind of normativity that methodological normativity is incompatible with is metanormativity of the ex cathedra style. Methodological normativity requires that a philosophical theory is empirically based and makes factual claims about actual scientific practice, which ex cathedra metanormativity rejects. By contrast, a metanormative theory that is sensitive to the empirical reality of scientific practice relies on methodological normativity because its normative claims must be informed by or connected in some other way to factual claims (see section 3.2). It is thus possible that a philosophical theory possesses all three kinds of normativity: a philosophical theory that makes factual as well as normative claims about epistemic or social norms in science (either by making claims about norms or by relating some nonnormative element to epistemic norms; see section 5.2) is methodologically metanormative and object normative.

One might wonder whether methodological normativity involves or depends on claims that are, themselves, metanormative or object normative. Typical methodological norms that implicitly guide theory development in the philosophy of science in practice make use of, for instance, specific notions of scientific success, and they rely on assumptions about what the goals of science are or should be. Hence, one could argue that methodological normativity presupposes metanormativity and object normativity. I think this claim is basically correct. Two points must be emphasized, however: first, this claim holds only for some methodological norms; second, this claim does not imply that methodological normativity reduces to metanormativity or to object normativity and can be eliminated from my typology of normativity in the philosophy of science.

Methodological normativity arises from normative claims that tell philosophers which types of empirical information are valuable (such as paradigmatic examples, instances of successful science, unbiased and clear-cut empirical information) and how they should proceed in developing a philosophical theory (for example, establish coherence, reveal conceptual connections, ignore biased or nonrepresentative cases). Methodological normativity is thus distinct from metanormativity because methodological norms tell philosophers what to value and what to do, whereas metanormativity concerns norms that tell scientists what to value and what to do (for example, what good cases of reductions are or which methods to apply). Methodological normativity is also distinct from object normativity because object normativity arises from norms in science, whereas methodological normativity traces back to norms in philosophy.

Despite the distinctness of all three kinds of normativity, some methodological norms presuppose assumptions about which epistemic norms are or should be accepted in science. An example is the methodological norm that a case of *E* is good and should be taken into account by philosophers because it plays a crucial role in achieving the aim of manipulating a specific disease. This methodological norm rests on the assumption that manipulation is accepted as a major epistemic norm in medical science. Other methodological norms presuppose assumptions about what should be regarded as scientific success, such as the methodological norm that empirical information about *E* from a certain scientific field is good because this field is especially successful in terms of making novel predictions. This methodological norm depends on the metanormative claim that the notion of success should be spelled out in terms of novel predictions. Hence, some methodologically normative claims rely on object-normative or metanormative claims, even though they are not themselves object normative or metanormative.

## 7. Conclusion

The goal of this paper is to draw attention to the fact that normativity in philosophy of science is a multifaceted phenomenon. Normativity is not a single feature that a philosophical theory either has or doesn't have. The paper articulates three different ways in which a philosophical theory about science can be normative. Each of the three kinds of normativity has a different origin. Methodological normativity emerges from norms involved in the philosophical methodology of developing factual claims about scientific practice. Object normativity is due to the normativity of the objects of philosophical theorizing. Metanormativity arises from the normativity of the claims that a philosophical theory contains.

Even though the practice turn in the philosophy of science is sometimes characterized as a "shift from normative to descriptive perspectives on science" (Soler et al. 2014, 15) my analysis shows that it is only a shift from metanormativity of the ex cathedra sort. The turn to scientific practice involves a move away from the approach of giving advice about what science really is and how to do proper science without considering the empirical reality of scientific practice. Other than ex cathedra styles, then, philosophy of science in practice is compatible with metanormativity. What is more, the practice turn is said to be a shift to more realistic views of science that also recognize the norms inherent in the pursuit of scientific knowledge and the deep intertwinement of science and society. Frequently, the practice turn will thus involve a turn to object normativity. Finally, even if philosophers of science in practice are not interested in making normative claims (metanormativity) or in reflecting on epistemic or social norms in science (object normativity), their theories about scientific practice will not be free from normativity, because their methodology inevitably involves normative assumptions about how to select, interpret, and evaluate empirical information, and how to mutually adjust philosophical theories and empirical information.

*Bielefeld University*
*Department of Philosophy*
*Postfach, 100131*
*33501 Bielefeld*
*Germany*
*kaiser.m@uni-bielefeld.de*

**Acknowledgments**

This project was partly funded by the German Research Foundation (FOR 1063). My special thanks go to Ken Waters, who discussed various versions of this paper with me. His ideas and suggestions always enhance my philosophical thinking. I would also like to thank my colleagues at the University of Cologne and at Bielefeld University for their helpful comments on previous versions of the paper.

**References**

Ankeny, R., H. Chang, M. Boumans, and M. Boon. 2011. "Introduction: Philosophy of Science in Practice." *European Journal for Philosophy of Science* 1:303–7.
Bechtel, W. 2008. "Mechanisms in Cognitive Psychology: What Are the Operations?" *Philosophy of Science* 75:983–94.
Bechtel, W., and R. C. Richardson. 2010. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Cambridge, Mass.: MIT Press.

Boumans, M., and S. Leonelli. 2013. "Introduction: On the Philosophy of Science in Practice." *Journal for General Philosophy of Science* 44, no. 2:259–61.

Brandon, R. N. 1996. *Concepts and Methods in Evolutionary Biology*. Cambridge: Cambridge University Press.

Carrier, M., D. Howard, and J. Kourany, eds. 2008. *The Challenge of the Social and the Pressure of Practice: Science and Values Revisited*. Pittsburgh, Penn.: University of Pittsburgh Press.

Chang, H. 2011. "The Philosophical Grammar of Scientific Practice." *International Studies in the Philosophy of Science* 25, no. 3:205–21.

Craver, C. F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford: Oxford University Press.

Elgin, C. Z. 1996. *Considered Judgment*. Princeton: Princeton University Press.

Fincham, J. R. S. 2000. "Reductionism Should Be Clarified, Not Dismissed." *Nature* 406:343.

Gallagher, R., and T. Appenzeller. 1999. "Beyond Reductionism." *Science* 284:79.

Giere, R. N. 1999. *Science Without Laws*. Chicago: University of Chicago Press.

Hempel, C. G. 1965. *Aspects of Scientific Explanation: And Other Essays in the Philosophy of Science*. New York: Free Press.

Hull, D. 1974. *The Philosophy of Biological Science*. Upper Saddle River, N.J.: Prentice-Hall.

Hüttemann, A., and A. C. Love. 2011. "Aspects of Reductive Explanation in Biological Science: Intrinsicality, Fundamentality, and Temporality." *British Journal for Philosophy of Science* 62, no. 3:519–49.

Kaiser, M. I. 2015. *Reductive Explanation in the Biological Sciences*. Cham: Springer.

Kinzel, K. 2015. "Narrative and Evidence: How Can Case Studies from the History of Science Support Claims in the Philosophy of Science?" *Studies in History and Philosophy of Science* 49:48–57.

Kitano, H. 2002. "Systems Biology: A Brief Overview." *Science* 295:1662–64.

Kitcher, P. 2011. *Science in a Democratic Society*. Amherst, N.Y.: Prometheus Press.

Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Lewis, D. 1999. *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.

Lloyd, E. A. 2006. *The Case of the Female Orgasm: Bias in the Science of Evolution*. Cambridge, Mass.: Harvard University Press.

Longino, H. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press.

Mansnerus, E., and S. Wagenknecht. 2015. "Feeling with the Organism: A Blueprint for an Empirical Philosophy of Science." In *Empirical Philosophy of Science: Introducing Qualitative Methods into Philosophy of Science*, edited by S. Wagenknecht, N. J. Nersessian, and H. Andersen, 37–61. Cham: Springer.
Mazzocchi, F. 2008. "Complexity in Biology: Exceeding the Limits of Reductionism and Determinism Using Complexity Theory." *EMBO Reports* 9:10–14.
McHugh, C. 2012. "The Truth Norm of Belief." *Pacific Philosophical Quarterly* 93:8–30.
McMullin, E. 1970. "The History and Philosophy of Science: A Taxonomy." In *Historical and Philosophical Perspectives on Science*, edited by R. H. Steuwer, 12–67. *Minnesota Studies in the Philosophy of Science 5*. Minneapolis, MN: University of Minnesota Press.
Nagel, E. 1961. *The Structure of Science: Problems in the Logic of Scientific Explanation*. London: Routledge.
Norton, J. 2003. "A Material Theory of Induction." *Philosophy of Science* 70, no. 4:647–70.
Paul, L. A. 2012. "Metaphysics as Modeling: The Handmaiden's Tale." *Philosophical Studies* 160:1–29.
Popper, K. 1959. *The Logic of Scientific Discovery*. Translation of Logik der Forschung. London: Hutchinson.
Rosenberg, A. 1985. *The Structure of Biological Science*. Cambridge: Cambridge University Press.
Ross, D., J. Ladyman, and H. Kincaid, eds. 2013. *Scientific Metaphysics*. Oxford: Oxford University Press.
Rouse, J. 2002. *How Scientific Practices Matters: Reclaiming Philosophical Naturalism*. Chicago: University of Chicago Press.
Sarkar, S. 1998. *Genetics and Reductionism*. Cambridge: Cambridge University Press.
Schaffner, K. F. 1974. "The Peripherality of Reductionism in the Development of Molecular Biology." *Journal of the History of Biology* 7, no. 1:111–39.
———. 1993. *Discovery and Explanation in Biology and Medicine*. Chicago: University of Chicago Press.
———. 2006. "Reduction: The Cheshire Cat Problem and a Return to the Roots." *Synthese* 151:377–402.
Schickore, J. 2011. "More Thoughts on HPS: Another 20 Years Later." *Perspectives on Science* 19, no. 4:453–81.
Schindler, S. 2013. "The Kuhnian Mode of HPS." *Synthese* 190, no. 18:4137–54.
Sober, E. 1975. *Simplicity*. Oxford: Clarendon Press.
———. 2008. *Evidence and Evolution: The Logic Behind the Science*. Cambridge: Cambridge University Press.

Soler, L., S. Zwart, M. Lynch, and V. Israel-Jost, eds. 2014. *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. New York: Routledge.
Thagard, P. 1988. *Computational Philosophy of Science*. Cambridge, MA: MIT Press.
van Thiel, G., and J. van Delden. 2010. "Reflective Equilibrium as a Normative Empirical Model." *Ethical Perspectives* 17:183–202.
Waters, C. K. 1990. "Why the Anti-Reductionist Consensus Won't Survive: The Case of Classical Mendelian Genetics." *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 125–139. Chicago: University of Chicago Press.
———. 2004. "What Concept Analysis in Philosophy of Science Should Be (and Why Competing Philosophical Analyses of Gene Concepts Cannot Be Tested by Polling Scientists)." *History and Philosophy of the Life Sciences* 26, no. 1:29–58.
———. 2008. "Beyond Theoretical Reduction and Layer-Cake Antireduction: How DNA Retooled Genetics and Transformed Biological Practice." In *The Oxford Handbook of Philosophy of Biology*, edited by M. Ruse, 238–62. Oxford: Oxford University Press.
———. 2014. "Shifting Attention from Theory to Practice in Philosophy of Biology." In *New Directions in the Philosophy of Science, vol. 5 of The Philosophy of Science in a European Perspective*, edited by M. C. Galavotti, D. Dieks, W. J. Gonzalez, S. Hartmann, T. Uebel, and M. Weber, 121–39. Cham: Springer.
White, R. 2005. "Why Favour Simplicity?" *Analysis*, 65, 205–10.
Wimsatt, W. C. 2006. "Reductionism and Its Heuristics: Making Methodological Reductionism Honest." *Synthese*, 151, no. 3:445–75.
———. 2007. *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge, Mass.: Harvard University Press.
Winther, R. G. 2011. "Part-Whole Science." *Synthese* 178:397–427.
Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.