1

PILLARS OF SCIENCE: REASONS, KNOWLEDGE, AND TRUTH

The empirical sciences are dedicated to describing, explaining, and predicting natural phenomena, such as ionic bonding, predation rates, formation of galaxies, bird migrations, gravitational waves, protein metabolism, or the probability of an earthquake in northern California. The ultimate goal is to gain *genuine knowledge* about such phenomena – what exactly they are, what regularities govern their occurrence, why and how they occur, when they might happen in the future, and with what probability.

In subsequent chapters, we will investigate many of the difficult and fascinating issues involved in a successful exploration of the natural world. Some of those issues arise with equal force in different sciences, while others may arise in only one discipline, but not in any of the others, at least not in the same form. For now, however, we are going to discuss some of the concepts and procedures that arguably can be found in all empirical sciences. These include (i) the fundamental nature of reasons we give for accepting or rejecting a scientific hypothesis, or any ordinary claim about the world for that matter; (ii) the general structure of the different kinds of inferences we can make on the basis of observations; (iii) the nature of truth and the specter of relativism; and (iv) the relation between facts, hypotheses, laws, and theories. Having a good grasp of these concepts will prove invaluable for understanding the intricacies of empirical science and the philosophical problems that arise in the course of scientific investigations.

The initial discussion will be a bit on the abstract side, and you might at first find it puzzling how questions such as the ones we are going to raise in

This is Philosophy of Science: An Introduction, First Edition. Franz-Peter Griesmaier and Jeffrey A. Lockwood.

this chapter are of any relevance to the practice of science. The situation is perhaps similar to learning the rules of baseball before swinging a bat. Sure, you might hit the ball if you don't know the how the game is played, but you might well run in the wrong direction afterward! And so, as our discussion proceeds in the later chapters, you will find that having gained mastery of some fundamental concepts in epistemology (the study of knowledge) is of great value. So let's get going.

1.1 Epistemic Reasons

Typically, when we wonder whether we should accept some claim, such as a scientific hypothesis, or not, we look for reasons for doing so. Another way of putting this is that we don't believe something without having reasons. For example, if someone asks you to accept that there are intelligent, extraterrestrial life forms, you're likely to ask for reasons before you adopt this belief. And if the other person just hems and haws, you're not going to believe in extraterrestrials. On the other hand, if you are presented with the cosmic background radiation as evidence for the occurrence of the so-called Big Bang, you have some reason for believing that the universe emerged through this sort of process. (Notice that when we use the term "belief," we do not mean to talk about religious faith.)

Of course, not any old reason is a good reason. If I believe something because I would like it if it were true, or because I am better off believing it, the belief might well turn out to be false – and in the overwhelming number of cases, it will be. Wishful thinking rarely leads to true beliefs. Thus, we need another kind of reason for believing something if we want to find out the truth about the world.

Reasons of the desired kind are called *epistemic* reasons. They are the sort of reasons that allow us to accept a belief only if there is good evidence for its truth, or only if the belief doesn't contradict other, already well-established beliefs derived from good evidence. Of course, it is very contentious what makes for a *good* epistemic reason. The debate over which (types of) epistemic reasons are to be preferred over others constitutes part of what's called *epistemology*, or the study of knowledge. Epistemic reasons are usually divided into two kinds: those which *guarantee*, in a sense to be specified momentarily, the truth of their target beliefs, and those that merely make the truth *more probable*. We start with the former.

1.1.1 Conclusive Reasons

The first type of epistemic reasons are called *conclusive* reasons. A reason (R) is conclusive for some belief (B) if and only if the belief *B must be true if R is true*. And this condition holds even if there is not just one reason for B, but also in cases in which B rests on many reasons. In more general terms, if all the reasons for a belief are true, and if they are conclusive reasons, then their target belief *must* be true. Conclusive reasons guarantee true beliefs, which is strongest basis one can have for believing something. So, how can we understand this definition?

A good example for conclusive reasons are the premises of a *deductively valid argument*. In such an argument, if all the premises are true, then the conclusion *must* be true as well. Here's a simple example:

Premise 1: All humans are mortal.
Premise 2: Stephen Hawking is human.
Conclusion: Thus, Stephen Hawking is mortal.

Clearly, if premises 1 and 2 are both true, then the conclusion is guaranteed to be true. Thus, the two premises together are conclusive reasons for believing that Hawking is mortal. But this sort of reasoning is not often helpful for advancing our scientific understanding of the world. Let's see why.

Notice that in a deductive argument, what's really happening is that information that is already contained implicitly in the premises, is made explicit in the conclusion. In other words, the conclusion does not reveal any new information. It restates the information that's already contained in the premises. That's why such inferences are *safe*: truth in – truth out.

Deductive reasoning (i.e., reasoning that proceeds by providing conclusive reasons) is mostly confined to two major disciplines: mathematics and logic. Yes, sometimes we use deductive reasoning in the empirical sciences, such as in cases in which we deduce observational consequences from a theory in order to test it (which can include refuting the theory):

Premise 1 (Theory 1)	All birds can fly.
Premise 2 (Theory 2)	Penguins are birds.
Premise 3 (Deduced Consequence)	Penguins can fly.
Premise 4 (Observation):	Penguins can't fly.
Conclusion:	Not all birds can fly.

However, a lot of scientific reasoning is nondeductive. Why? Because typically, in scientific reasoning, we want to infer something about the world at large on the basis of a limited number of observations. Such inferences are inherently risky because their conclusions convey information that goes beyond the information contained in the descriptions of the actual, limited observations that have been made.

For example, if I infer, on the basis of having observed the eating habits of 20 koalas, that all koalas eat eucalyptus leaves, I make such a risky inference. I assume, among other things, that the koalas I observed are typical of their species. This assumption could easily be wrong, as I might have come across a peculiar band of koalas that happen to consume eucalyptus. That such inferences are risky, however, doesn't show that they are altogether unreasonable. The conditions under which they are reasonable are somewhat difficult to pin down, and we will tackle this challenge in the next section.

Now, given that reasoning nondeductively is risky, and that the conditions of its reasonableness are somewhat elusive, one might think that science should aim at just using deductive inferences, precisely because they are safe – even certain. But that would be a mistake. Remember: They are safe because in an important sense, they are uninformative. Since there is no new information in a deductive conclusion that was not already implicitly contained in the premises, deductive inferences won't allow you to gain more information about the world by reasoning from your evidence. To accomplish this, we need to go beyond an obsession with certainty, which is provided by conclusive reasons and reasoning, and enlarge our toolbox. The tools we need, especially for the empirical sciences, are various forms of *defeasible* reasoning, and thus *defeasible reasons*.

1.1.2 Defeasible Reasons

The second, and much more common, type of epistemic reasons are called *defeasible* reasons. They are also sometimes called *probable*, or *prima facie* reasons. The main difference between these and conclusive reasons is that even true defeasible reasons don't guarantee the truth of their target belief. Consider this example:

> You are near the mouth of a cave looking at a rock formation just inside the cave. The formation looks red to you. This "red-looking" is a good (defeasible) reason for believing that the rocks *are* red. However, as we all know, lighting conditions vary in natural settings and can be deceptive. Thus, it could be the case that the rock formation isn't really red; its red appearance could be produced by weird lighting filtering into the cave. Thus, although the red-appearance of the formation is a good defeasible reason for believing it to be red, the truth of this latter belief is not guaranteed.

What defeasible reasons do is to make the truth of the belief for which they are reasons *probable*. (That's why they are also called probable reasons.) A red appearance of a rock formation makes it more probable that the formation is red than that it is not. Of course, you could acquire a further bit of information which *defeats* the strength of the reason (that's why they are called defeasible reasons). For example, you could notice that there is a brilliant sunset outside the cave, which makes it likely that many even nonred things look red. Thus, the fact that the formation looks red to you is no longer a very good reason for believing that it is actually red, given that many nonred things will seem to be red in these lighting conditions.

There are actually two recognized kinds of defeating information, or defeaters: so-called *rebutting defeaters* and *undercutting defeaters*. In the example just given, noticing that there is a sunset is an *undercutting* defeater. It undercuts the evidential force of your original reason for believing that the formation is red. Given that you know of the red lighting, you now can't fully trust that things have the color which they seem to have. Of course, the rocks could still be red. But you would need to illuminate them with a white light or take a sample and observe it during daylight to make sure. On the other hand, it could also be the case that a geologist tells you that the formation isn't red, because there are no rocks of this color in the region. To the extent that you can trust her, you now have a *rebutting* defeater for your belief that the formation is red.

As mentioned above, defeasible reasons constitute the vast majority of the reasons we have for believing something. Conclusive reasons are limited to mathematics and logic. It is therefore extremely important to remember that talk of "physical proof," for example, is a misunderstanding of the concept, if "proof" is being used with its technical meaning. A proof consists in providing conclusive reasons for some target belief. That means that it must be literally impossible for the premises (i.e., reasons) to be true *and* the target belief to be false at the same time. Such high standards of evidence are unavailable in the empirical sciences. We just can't be certain. Even the best empirical evidence, on the basis of which we can form true premises for an argument, does not guarantee the truth of the target belief (i.e., hypothesis or theory; we'll explore the difference between these below).

Good evidence makes the supported theory highly probable, often so probable in fact that we may legitimately say that we justifiably assert that it is true. But the truth is not guaranteed by even the best evidence. Thus, there can't be literally a proof of any empirical belief – not the Big Bang, quantum mechanics, plate tectonics, or evolution. Proof requires conclusive reasons which are limited to mathematics and logic. This seemingly technical point is borne out by the history of science. Most, if not all, empirical theories of the past have turned out to be strictly speaking false, even though many were in the right ballpark and much better than the alternatives. Had the evidence used to support them guaranteed their truth, they couldn't have turned out false. But they did. Let's look at the infamous "substance" *caloric*, the stuff heat is allegedly made of – a once entirely reasonable and well-supported hypothesis that was ultimately wrong.

There was once good evidence for the theory that heat is a sort of subtle fluid, exchanged between bodies. This fluid was called "caloric." We still use the term "calories," which derives from that old theory of heat. It was partially based on the observation that if two objects at different temperatures are in contact, they will eventually reach thermal equilibrium: the hotter object will cool off and the cooler object will heat up, until they both have the same temperature. This looks a lot like what happens when one opens the valve of a hose connecting two buckets that hold different amounts of water. The water will run from the bucket with more water to the other one, until both hold the same amount of water. If heat was also a fluid, we would have a good explanation of the fact that adjacent bodies eventually reach thermal equilibrium. So based on this evidence and reasoning from what seemed to be an analogous system, scientists inferred that heat is such a fluid and called it caloric.

Of course, now we know that heat is, roughly speaking, micromotion or the movement of unobservably small particles (atoms and molecules). This understanding arose from the production of canons in the eighteenth century.[1] Using a cold drill to bore a hole into cold metal produced a lot of heat. But where was the heat coming from? Caloric was assumed to be a conserved quantity that could only be redistributed among objects, but neither created nor destroyed. The "new" heat from the canons was tantamount to running a hose between two half-full buckets and having them both overflow. You might think that this problem would've been apparent from other instances of friction, but the phenomenon was dramatically evident in the manufacturing of canons. What this episode shows is that even good evidence for a theory provides at best a defeasible reason for accepting that theory. New evidence can force us to retract our theory, exactly like knowledge of weird lighting conditions in a cave can force us to retract our belief about the color of rock formations.

1.2 Reasoning from Evidence

Scientists most often make use of *induction*, or drawing conclusions from evidence. Generally speaking, induction is *an inference from the observed to gain information about the unobserved (or unexamined)*, and it takes three different forms: *statistical inference, inductive generalization*, and *inference to the best explanation* (*IBE*).

It is universally recognized that what these three forms of reasoning have in common is that they are defeasible. Beyond that, the terminology here is unfortunately not as widely agreed upon as it is with respect to deductive reasoning, but we can at least try to distinguish clearly between the three different forms of inductive reasoning just mentioned. Notice though that some textbooks in the sciences define inductive inferences as inferences from the particular to the general. This is misleading, because it covers only a tiny fraction of inductive inferences. We start with statistical inference.

1.2.1 Statistical Inference (SI)

The simplest example of an inductive inference is that of inferring something about an entire population from observing only some of its members. Recall the earlier example involving koalas: I inferred from having observed 20 of them munch exclusively on eucalyptus leaves that *all* members of the species *Phascolarctos cinereus* (that's the koala's scientific name) feed on eucalyptus leaves. Of course, such inferences are not restricted to biological populations. We might conclude that all igneous rocks are black, after we have seen many lava fields and observed that all of those were black. We can characterize the nature of statistical inferences in the following way:

> A statistical inference is an inference from the observed frequency of a property in a sample to the claim that the same frequency holds for the population from which the sample was taken, within a certain margin of error.

Here is an example in explicit form:

Premise 1:	The frequency of red marbles in a sample of 200 balls drawn from an urn was 49%.
Premise 2:	The urn contains exactly 1,000 marbles which are either red or black.
Conclusion:	The frequency of red balls in the urn is 50%, with a margin of error of ± 2%.

Obviously, SI is an inference from the observed (the sample) to the unobserved (the population). Suppose you randomly picked up the first one hundred plants in a meadow and every one of them was a grass. You might well infer that every plant in the field was a grass. As we all know, beliefs (or hypotheses) based on SI can turn out false. Not all igneous rocks are black, and it's unlikely that all plants in a meadow are grasses, although koalas seem to invariably eat eucalyptus. Often, this is due to sampling problems, which can never be fully eliminated (maybe all the tall plants that are easily accessed are grasses, but some small, ground-hugging plants are broad-leaved species). But even if the sampling doesn't involve any bias, evidence from samples provides only defeasible reasons for beliefs about the relevant population, as the deviation of election results from predictions based on sampling (called polling) clearly demonstrates. There is much more to be said about SI, some of which you'll find in later chapters.

1.2.2 Inductive Generalization (IG)

This form of inference is a bit more difficult to characterize to any great degree of precision. In fact, not even the name is widely agreed upon. Sometimes, IG is used to refer to what we call SI. Since nomenclature is a matter of convention, nothing really turns on it, as long as we are reasonably clear about the differences among the kinds of inferences. In order to begin developing a good understanding for what we decided to call IG, it's best to start with an example.

Suppose you are interested in determining the functional relation between the period of a pendulum (how long it takes to pass through one cycle) and the length of its string. Dutifully, you plot changes in the dependent variable (the period) against variations in the independent variable (the length of the string). Unavoidably, you'll get a general trend with a somewhat messy point distribution. If you were to precisely connect all the points, you'd end up with a jittery line. "Nature can't be that crazy," you mutter to yourself, as you begin accepting that some of the points might not fall exactly on the line describing the actual relationship. You know about air resistance, the variable elasticity of the string due to changes in ambient humidity levels, the imprecision of your starting and stopping the timing device, and other factors that really have nothing to do with the true relation between length and period (that's why we call those factors "noise"). Thus, you decide to go for a nice, neat line – a section of a parabola, as it were. Then, you find the algebraic expression that generates that line. Finally, you make an inductive generalization and conclude that the period T of all pendula is related to the length l of their respective strings as follows: $T=2\pi\sqrt{l/g}$, where g is the gravitational acceleration. In fact, you are proud to have discovered the ideal pendulum *law*, which holds for all pendula with sufficiently small angular displacement. We will say more about the question of what a law of nature is in Chapter 12.

1.2.3 Inference to the Best Explanation (IBE)

This form of inference is exemplified by the story of caloric. The idea is that we observe some regularity and then postulate one or more mechanisms (or entities) that could be responsible for the observed regularity. If we can come up with just one, and it strikes us a plausible, we call it a day and accept it (this simplification will be corrected momentarily). If there is more than one, we look for the explanation which strikes us as best. To see what's going on here, look at the following example, inspired by philosopher Elliott Sober.[2]

Suppose you are sitting in your living room and suddenly hear strange noises coming from your attic – a quick succession of what sounds like little taps and then a rumbling noise. You consider two hypotheses: First, the noise is produced by gremlins from outer space that have landed on your roof and are now bowling in your attic. Call this hypothesis G. Second, the noise is produced by the neighbor's cat, which got into your attic and is trying to catch mice, but keeps running into the books you have stacked up there. Call this hypothesis C. Clearly, C is better than G, although G is an explanation of sorts: If it were true, then the probability of hearing those noises would be quite high. However, the same is true of C, and since C is more plausible than G in light of all the other things you believe, C is clearly the better explanation. You then infer that the best explanation is the most probable one and thus accept C. You have inferred the best of the explanations under consideration; this is IBE.

Several things are worth commenting on. First, both G and C make the observation probable – both bowling gremlins and mouse-chasing cats could produce those noises you hear. Second, both C and G might be false. Maybe it's neither cats nor gremlins, but it's some neighborhood kids playing a practical joke. Third, and related to the last point, what hypothesis counts as the best is partially determined by which ones you can come up with. In other words, the best hypothesis from among those we thought of need not be a very good one, all things considered. Often, we miss an even better hypothesis, as we know from the history of science and discuss further in Chapter 12. In fact, the history of scientific progress is one of not only gathering more evidence but improving explanations such that was once the "best" is supplanted by something better.

Thus, even if you are very creative in generating hypotheses, you might generate really awful ones and shouldn't believe any one of those. This has prompted some to eschew the use of IBE altogether, especially insofar as it pertains to unobservables (things we can't directly see or otherwise sense, such as electrons or magnetic forces). In short, IBE can provide some reason for accepting a claim (we use it in forensic sciences all the time, for example, when we try to find the person whose presence is the best explanation of all the clues, and infer that the person who best fits the clues is the perpetrator), but it certainly doesn't guarantee our knowing the truth.

Finally, it is important to point out that IBE cannot be reduced to other forms of inductive inference. Inferring the presence of a stray cat in my attic as the best explanation of the noise I am hearing does not (need to) involve prior observations of stray cats *in my attic* and their behavior. Thus, this inference is different from a statistical inference, which, if you recall the example of koalas, does rely on observations of the feeding habits of a number of koalas to infer something about what other koalas will eat. Neither am I trying to establish any sort of regularity when I infer that a cat must have gotten into my attic.

I am simply interested in explaining this particular and odd event by evaluating various hypotheses as to their plausibility in light of my background knowledge about cats, none of which I need to have oberved in an attic.

Let's recap what we discussed in this section. We distinguished between conclusive and defeasible reasons. Along with that distinction comes another one, namely, the distinction between deductive and inductive reasoning. The former can be characterized as an inference from the implicit to the explicit, while the latter is an inference from the observed to the unobserved. The following diagram (Figure 1.1) shows the distinctions you should keep in mind:[3]

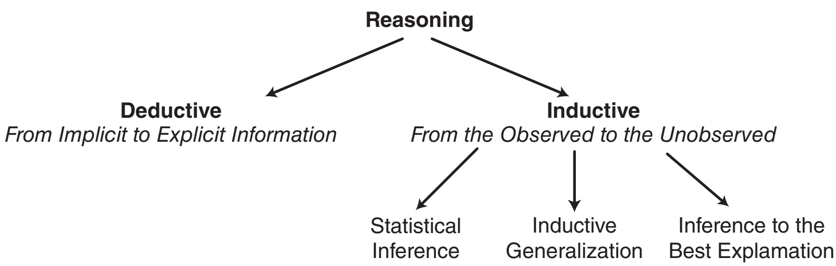

Figure 1.1 Forms of reasoning.

1.3 Knowledge and Truth

Scientific inquiry is often characterized as an especially promising way to increase our knowledge of the world by discovering more and more truths about it. But what exactly is knowledge, and what is truth? To answer the first question, we can start by determining what the difference between mere belief and knowledge is. Suppose Marvin strongly believes that atoms are held together with hooks. Does Marvin know this to be the case? It seems not, for the simple reason that it is false. We can't know what's false. Of course, we can believe what's false, and we can falsely believe that we know something even though it is in fact false. But we can't know that 2 + 3 = 78. Thus, we can conclude that in contrast to mere belief, a belief that counts as knowledge has to be true.

But being true is not enough for a belief to count as knowledge. Suppose I am terrified of snakes, flipping a coin every morning before entering the lab to see whether I have to use protective gear. This morning, the coin landed heads up, which means, to me, that there is a snake in the lab. Moreover, there in fact is a snake – unnoticed by me, a practical joker snuck it into the building. Do I know that there is a snake? It seems not. That the coin landed heads up doesn't provide a good reason for believing that there is a snake. And yet it is true. This shows that for knowledge, we don't just need truth – we need good reasons as well.

Putting it all together, knowledge amounts to true beliefs that are based on good epistemic reasons. Some of the readers might be aware of the fact that over the last several decades, it has become clear that even true beliefs for which I have good reasons might not automatically qualify as knowledge – something else is needed, at least in the kind of instances usually called *Gettier* cases.[4] However, because we encounter such cases only rarely, if at all, during scientific research, we won't discuss them here. In the next couple of chapters, we will instead explore the notion of evidence and how it is supposed to yield good reasons for our beliefs about the world, thus providing us with one ingredient of scientific knowledge.

The second ingredient is truth. There are many competing theories of what truth is, and there is an important debate in the philosophy of science whether or not our best theories are true. We can't explore the first question here, but we will explore the debate about the alleged truth of scientific theories in Chapter 12. For now, we can just stipulate that by a true theory, or a true belief for that matter, we mean a theory (or a belief) that corresponds to the facts in the world. For example, the hypothesis that you are reading this sentence right now is true because it corresponds to the fact that you are reading it.

1.4 Facts, Hypotheses, and Theories

This may be a good point for some remarks about the relation between facts, hypotheses, and theories. To begin with, facts are actual states of affairs in the world, or how things are. For example, there is the fact that you are reading this sentence right now. This is something that is happening in the world – you reading this sentence is among all the facts that together determine the world. In contrast, hypotheses and theories are (often) linguistic expressions of our beliefs: They consist of sentences purporting to report the facts. For example, the sentence "You are reading this chapter right now" is a hypothesis about you and what you are doing. The hypothesis is right now true (you are still reading, aren't you?), but once you get bored and look for your friends on social media, it will become false. And this happens because the facts will have changed.

The difference between fact and either theory or hypothesis can be easily overlooked. For example, in discussions about evolutionary biology and creationism, one often hears creationists complain that *evolution is just a hypothesis*, while the evolutionists call the *theory of evolution* a *proven fact*. Both sides here are equally guilty of speaking carelessly. First, *evolution* is not a hypothesis, because evolution is not a linguistic entity. It is a process that occurs in nature. *Evolutionary theory*, on the other hand, is a linguistic entity (it's a bunch of sentences written in books), which can be either true or false. And, of course, evolutionary theory is true if and only if (iff) it corresponds to the facts that comprise the process of evolution, such as the facts concerning natural selection.

Second, there are no such things as "proven facts." As we have already seen, only mathematical or logical sentences (and the propositions expressed by them) can be genuinely *proven*. Empirical theories can only be justified by defeasible reasons, which, however, are often strong enough to warrant treating those theories as reporting the facts correctly. Moreover, facts are not the sort of thing that can be proven. You can't deduce a fact from some premises, because they are not linguistic entities. What you can do is try to bring about some fact, by acting in various ways (e.g., you could apply antibiotics to bacteria to bring about the evolution of resistance). Or you can discover facts. But you can't prove them – they are the wrong kind of thing for that.

Finally, what's the relation between hypotheses and theories? It seems that we often try to use the distinction to mark a difference in how well supported an empirical claim is and how comprehensive its content is. We demand better support for a theory than may be required for a (mere) hypothesis, which is often thought of as something like an educated guess, a conjecture, or a sort of hunch.

If we approach the matter a bit more systematically, we can think of hypotheses as relatively isolated or "stand alone" empirical claims. For example, in scientific modelling, we are often only interested in correctly identifying relations between dependent and independent variables (such as period and length in a pendulum). This can be done regardless of how the resulting model fits together with other things we believe.

When we propose a theory, however, we typically try to fit several hypotheses together into a coherent whole. Recall the earlier example involving the ideal pendulum law. During the process of fitting a curve to the data points you collected, you might think of the result as your hypothesis about how period and length are related. You can then go on and embed this hypothesis into a comprehensive theory or encompassing model, such as the theory of simple harmonic oscillators, or even more broadly, into all of classical mechanics.

"Embedding" means that you try to show how your hypothesis can be derived as a special case of a simple harmonic oscillator or as an instance of Newton's laws of motion (together with some boundary conditions). Such an embedding also means that the evidence for the other parts of the theory becomes indirect evidence for your original model.

Thus, we can draw the distinction between hypothesis and theory in terms of how many different pieces of information have been integrated, although we have to keep in mind that this distinction is vague.

1.4.1 "It's True for You but Not for Me"

One topic that often comes up in discussions about truth – a topic we haven't touched on yet – is that of *relativism*. Many people profess to be relativists about truth. They might say things such as: "This is true for you but not for me," or "It was true for Ptolemy that the earth is stationary, but now it is no longer true." Some might even provide a theoretical defense of the view that truth is relative to one's culture.

What is often in the background here is the idea that one needs to respect another person's opinion and perspective. After all, wars have been fought over differences in beliefs. So epistemic humility is a virtue. But this gives rise to a problem. If I disagree with another person, I seemingly need to hold that truth is a matter of opinion, perspective, or culture. However, there is an alternative.

One can respect other people's opinion and perspective in that one can appreciate how their evidence and cultural experiences might make it seem that so-and-so is the case, while my evidence makes it seem that this is not so. Remember, the defeasibility of our reasons is a sound basis for doubt. The evidence people have access to often varies dramatically among individuals.

It might even be the case that individuals have access to the same body of evidence but give different weights to the various elements based on experience, existing beliefs, practical interests, and cultural values. We'll explore this possibility in some detail in the next chapter, but for now let's just say that once we establish the spatiotemporal context (and sometimes the values), we often find that there is little disagreement. Once again, in philosophy as in science, one must take care to make matters clear and, in this case, to not mistake differing interests or perspectives for there being relativistic truth.

Thus, what it is reasonable to believe can vary dramatically from person to person. We have already seen earlier that it can be quite reasonable to believe something that's false. The evidence one has might in fact not be very good evidence. Thus, I can respect opinions that are different from mine for their reasonableness without having to claim that truth varies from person to person. Once we make a distinction between two ways in which beliefs can be "good" (i.e., their reasonableness and their truth), we can explain variations among what reasonable people believe in terms of differences in evidence without having to be relativists about truth.

And that is a good thing, because relativism about truth is ultimately indefensible. Take, for example, the sentence, "There is no objective truth." Is that sentence itself true? If it is objectively true, then it has just refuted itself. If it is merely relativistically true, then why should I, who believes in objective truth, be moved by it? Subscribing to relativism about truth puts one into a really bad intellectual (and practical!) position.

Of course, these remarks are far from sufficient for putting relativism to rest. There are fairly sophisticated defenses of this position, but it would lead us too far into epistemology proper to give them a fair hearing. What the remarks do show, however, is that relativism as a knee-jerk reaction to the existence of disagreements about truth is simply too naïve.

1.4.2 Perspectivism

Although relativism, at least in its unsophisticated form, is not a plausible response to disagreements about what is true, there is a viable approach that allows for the truth of multiple claims about the world, but not the truth of every claim as relativists would contend. Ronald Giere advocated for *scientific perspectivism*, a philosophical position mediating between the hard-driving objectivism of most scientists (as well as the unyielding realism of many philosophers of science), and the "anything goes" constructivism found among postmodern social scientists and humanists. Perspectivism contends that while there is an objective reality, our statements about its features are always made from a particular viewpoint (both literally and figuratively) and are hence partial truths, at best. This position fundamentally differs from relativism in that once a perspective is chosen, claims about the natural world can be empirically tested and refuted or confirmed.

The notion that truth claims are always relative to a perspective is, as Giere points out, not radical. Even the staunchest empiricists cannot escape the fact that scientific claims are relative to language. What Giere initially proposed was to extend linguistic contextualization to one's physical position (both in relation to the object of study and the internal processing of the incoming signal).

Consider, for example, a biologist who asserts that the petals of a particular flower are white. One need only look to see that this is true. However, suppose another biologist comes along with a device which reveals ultraviolet reflectance, looks through this instrument at the flower, and declares that it is striped. And, in fact, some flowers have so-called "bee guides" that direct bees, which can see in the ultraviolet, to the stamen and nectary which provide pollen and nectar. Which description of the flower is true? The perspectivist would respond that the truth of the competing claims depends on whether the accounts are rendered from the viewpoint of a human or a bee. Note however, that once the perspective is fixed, the descriptions can be wrong; from the human perspective it would be mistaken to call the flower red, and from the bee perspective it would be wrong to assert the flower is checkered.

While this example illustrates that scientific claims are conditional on biophysical and instrumental perspectives, Giere also maintained that propositions are made within a theoretical context. For example, a sociobiologist might contend that helping a neighbor is explained in evolutionary terms by reciprocal altruism, while a cultural anthropologist might argue that cooperation serves to strengthen social bonds in a community. Along similar lines, a neurophysiologist might explain that people eat sweet foods because sugars stimulate nerves leading to the gustatory and reward regions of the brain, while an evolutionary biologist might explain that sweetness is favored by its association with concentrated sources of energy which were rare in the history of our species. Giere maintained that extreme pluralism arises from a multitude of competing, theoretical perspectives in empirically weak scientific fields at the limits of experimental capabilities (e.g., gravity waves).

Others philosophers have pushed perspectivism further, to take into account the possibility of different interests. That is, to assess a scientific claim, the needs and desires of the investigator are relevant. When asked if a forest fire should be allowed to burn, an ecologist might answer that the dominant trees in the ecosystem of concern are pyrophytic, requiring fire to reproduce, so "let it burn." But another ecologist might answer that the last surviving population of the endangered yellow-spotted salamander lives in the ponds within the burning forest and the fire should be extinguished to conserve this species. And an ecologist working for a forest products company might insist that jobs would be lost if the fire is not controlled.

You can see that perspectivism admits that scientists may have differing values and, in fact, there is no value-free science (at a minimum, scientists value knowledge). If the forest fire is allowed to burn it is true that the plants will recover (a long-term perspective valuing natural processes), that the amphibian will go extinct (a perspective valuing the existence of species), and that people will be harmed (a short-term perspective valuing economic well-being). So, while logically incommensurable claims are not countenanced from a given perspective (e.g., it cannot be the case that the salamander is both endangered and flourishing at this time), incommensurable conclusions can be drawn from scientific data when perspectives diverge (e.g., forest health might be harmed in the short-term and benefited in the long-term by the same event).

1.5 Conclusion

This completes our brief survey of some central concepts in epistemology, such as the concepts of epistemic reasons, knowledge, and truth. As you can imagine, there is much more to be said on all of the topics we touched on. The suggested readings provide further and deeper treatments of many of those topics, if you are interested. Having shown that scientific knowledge most often involves some form of inductive reasoning from evidence, we next go one step deeper and turn to a fundamental question in the philosophy of science: What is the nature of evidence, and how can we access evidence?

Notes

1 The debate between a dynamic theory of heat and the defenders of the caloric account was an important stage in the development of thermodynamics. Early defenders of the dynamic theory included Galilei Galileo, Isaac Newton, and Robert Boyle, who thought that heat is due to the internal motion of particles. This however was difficult to reconcile with the discovery of latent heat (i.e., the heat absorbed by a substance that undergoes phase change, e.g., melting, but doesn't rise in temperature) in 1757 by Joseph Black. For details, see Robert D. Purrington, Physics in the Nineteenth Century, Rutgers University Press, 1997, pp. 75–101.

2 Elliott Sober, Philosophy of Biology. Westview Press, 1993, p. 32.
3 A form of inference related to inductive inferences is reasoning by analogy. We'll say more about this in Chapter 9.
4 A very simple Gettier case involves a clock whose hand is stuck at 5. If you don't know the clock is broken, you are justified in believing the time it indicates (it has worked reliably so far). Moreover, if you just happen to look at that clock when in fact it is 5 o'clock, you have a true belief that is also justified. But few would say that this justified true belief counts as knowledge of the time, because your belief is only accidentally true.

Annotated Bibliography

Ronald Giere, 2006, *Scientific Perspectivism*. Chicago: University of Chicago Press.

Giere argues that scientific claims are conditioned on highly confirmed theory or reliable observations such that if the assumptions of the underlying theories or the biophysical properties of the instrument/observer changes, then the result can be a different, even incommensurable, accounts of the natural world. There is no complete, objective, or singularly correct account of the world through correspondence how the world really is.

Alvin I. Goldman and Matthew McGrath, 2015, *Epistemology. A Contemporary Introduction*. New York and Oxford: Oxford University Press.

This book contains advanced discussions of the structure and nature of justification, the Gettier problem, skepticism, and contextualism. It also covers methodological questions concerning the use of intuitions as evidence as well as important new developments in social epistemology which are important for understanding how knowledge emerges from research teams, as opposed to individual scientists.

David Manley, 2019, *Reason Better: An Interdisciplinary Guide to Critical Thinking* (Tophat Monocle) https://app.tophat.com/e/455176/assigned

This recent and very accessible book on critical thinking covers many of the concepts we encountered in this chapter, such as reasoning and inference. In addition, it contains illuminating discussions of various biases, inferential fallacies, and cognitive illusions that can interfere with successful reasoning.