Philosophy and Ethics of Science

9th Grade Mini-Course

Penn Project for Philosophy for the Young

Instructors/Contributors: Nicholas Friedman, Stephen Esser, Clarissa Busch, Lillian Phung

Contents

I. Introduction

In 2018 we launched a course called "The Philosophy and Ethics of Science" for 9th graders at the Science Leadership Academy (SLA), a high school in the Philadelphia School District. SLA offers students a choice of "mini-course" electives in non-traditional subjects, and our offering was developed as one of the options (students are graded based on attendance and participation). The class has usually consisted of seven or eight 90-minute sessions each school quarter (the lesson plans below include suggestions for adapting to shorter class times).

SLA has a science and technology focus, and our course complements this by using scientific themes as a launching point for engagement with philosophy. Our objectives are threefold: to introduce the students to philosophical topics and cultivate their reasoning skills; to have them gain experience engaging with challenging ethical issues; and to deepen their understanding of science by reflecting on its complexities. We also hope to motivate the students to further explore these ideas in their other coursework and in their everyday lives as consumers of scientific information. The sessions are organized by topic, with added modules on reasoning and argumentation and one or two classes devoted to ethics case studies. The class activities feature small group work, debates, and role-playing (a tabletop exercise and a mock trial). The emphasis is on independent thinking, constructive dialogue, and respect for other viewpoints.

We have focused on the following themes and key questions. *The Demarcation Problem*: How do we distinguish science from non-science and pseudo-science? *The Nature of Scientific Evidence*: What counts as evidence? What makes some evidence preferred over others? *Scientific Reasoning*: What kind of reasoning supports scientific claims? What are facts vs. opinions? Do scientists prove claims? *Science and Moral Questions*: Can science answer moral or ethical questions (or contribute to their answer)? *Values and Objectivity*: How do human values influence the sciences? Should we aspire to a value-free ideal? *Ethics Case Studies*: What ethical problems arise in research? What moral principles guide application of scientific discoveries?

In reflecting on this initiative, we feel that these 9th grade students were able to critically reflect on their intuitions and assumptions regarding science. Many of the students have clearly enjoyed getting the opportunity to discuss and debate deep questions. At the end of each course we survey the students, and the responses have been positive on format and content, with criticisms focused on a desire for even more interactive (and “fun” or “exciting”) elements. Challenges include the variation in the background knowledge of students, and keeping everyone on pace to complete activities. We intend to continue developing and improving this course while looking for other applications for this or similar curricula. Below we present eight lesson plans followed by some additional discussion of the themes of the course (including suggestions for further reading).

II. Links to Lesson Plans

Note: Lecture slides and other supplemental materials available upon request.

Lesson Plan 1: Course Introduction & Natural Disaster Response Tabletop Exercise

Lesson Plan 2: The Demarcation Problem

Lesson Plan 3: The Nature of Scientific Evidence

Lesson Plan 4: Arguments and Scientific Inference

Lesson Plan 5: Science and Morality

Lesson Plan 6: Values and Objectivity in Science

Lesson Plan 7: Ethical Case Studies

Lesson 8: Pharmaceutical Fraud Mock Trial

III. Philosophical Themes for Lesson Plans

A. Overview

As an academic discipline, Philosophy (from Greek, philosophia, "love of wisdom") is the study of fundamental questions about the nature of the world, people, and our relationship to the world and to each other. It encompasses subfields such as the following:

- Metaphysics: Study of the nature of reality and existence of things.
- Epistemology: Study of the nature and extent of knowledge.
- Value Theory: Study of principles and value. This can be broken down into sub-fields such as ethics: (study of morality and how we ought to live and act) and aesthetics (study of the nature of beauty and value in the arts).
- Logic: Study of reasoning and argumentation.

Of course, many questions in philosophy combine multiple branches. Although these fields are broad topics on their own, they can also be used in an applied way in order to examine other topics or fields. Examples include such disciplines as the history of philosophy, political philosophy, philosophy of mind, and philosophy of language.

Philosophy of science uses these approaches to study the nature and methods of science. The themes discussed in subsequent sections below represent just a subset of the topics explored in this discipline. The course title reflects the special emphasis on the intersection of ethics and values with science.

While philosophy and its sub-fields can be thus described in terms of its subject matter, it is perhaps more important to introduce philosophy to students as an activity or practice of thought.

Philosophy is a specific type of thinking that is:

- Reflective/Curious: Philosophers reflect on fundamental and basic questions regarding the nature of the world, their experience, and the lives of others. Philosophers are deeply curious and are willing to explore almost any question.

- Critical/Skeptical: Philosophers critically evaluate claims for support and evidence, and refrain from making judgements based on a hunch or inclination. Philosophers are skeptical of unsupported claims. Philosophers are similarly willing to accept that their own beliefs may be incorrect.
- Creative: Philosophers create new ways of thinking about the world and human experience.
- Rigorous: With a focus on precision, argumentation, and logic, philosophy is rigorous. Philosophy is a challenging discipline that requires the utmost intellectual virtues.
- Collaborative: Although philosophers are independent thinkers, philosophers work collaboratively. Philosophers build upon the work of others to make progress over time.

For the purposes of a philosophy class, this last point about collaboration is of greatest importance. Dialogue and discussion are crucial: we do philosophy together, allowing us to receive immediate feedback and build upon the ideas of others. However, care must be taken to ensure that conversations are productive: i.e., that they result in the development of new ideas, clarification of concepts, and generation of arguments in support of (or critiques of) well-defined, focused positions. Ideally, they result in individuals revising or altering their stance – as opposed to merely the back-and-forth sharing of claims in which individuals are already committed. Similarly, care must be taken to ensure that conversations remain civil, respectful, and inclusive. The instructor/facilitator must ensure that all voices are encouraged and heard, and that students are genuinely working to support and build upon the ideas of others – rather than allowing discussions to serve as a competition for who can best advance their own positions.

Further reading:

Andrea Borghini: “What is Philosophy” https://www.thoughtco.com/what-is-philosophy-2670737

B. The Demarcation Problem

What is science? What makes it different from other disciplines? What gives it a special role in our pursuit of explaining phenomena, making predictions, and developing technologies? While we can point to the subject matter and methods of science as they have been conventionally defined, philosophers of science think that a way to deepen our understanding of science is to try to find criteria for distinguishing it from subjects/activities that are not science. This exercise also becomes of practical importance due the need to identify practices that try to illegitimately claim to be science (pseudosciences).

It turns out to be very difficult to find a single criterion that defines an exact boundary between science and non-science. However, there are some key characteristics that usually characterize science, and the good news is that these do effectively help us recognize pseudo-science.

A central idea is that scientific claims/hypotheses should be tested/supported by evidence. But justifying claims is important outside of science, too. Other possible criteria likewise fail to draw perfect boundaries. Examples include asserting that scientific claims should be confirmed by publicly available observations, or that claims should potentially be falsifiable by experiment. In the first case, philosophers came to conclude that many of the accepted findings of science cannot be straightforwardly reduced to facts about simple observations. In the case of falsifiability, the criterion fails to accurately describe much of scientific activity. While claims need to be tested with available evidence, scientists do often strive to support claims rather than directly seeking to falsify them. Also, it is important not to put undue stress upon experimentation in particular, since this tends to exclude some sciences (like astronomy and paleontology) where experiments are not practical. More importantly many pseudo-scientific claims are also indeed falsifiable in principle -- it is just that their advocates tend not to acknowledge contrary evidence.

This observation about pseudoscience serves to point us in a different direction: to many of the norms and practices that characterize science. These include transparency of methods and findings, replicability of results, and institutional checks and balances such as peer review. Science succeeds by learning from its failures and because it has institutional structures that work to limit both inadvertent mistakes and fraud (albeit imperfectly). These features are what allow science to make progress, while certain pseudo-sciences (astrology, say) look little different than they did centuries ago.

A better approach is to come up with a list of criteria, while acknowledging that no sub-set of them will draw a perfect boundary in all cases. Here are some broad features that can help with the problem of distinguishing science from non-science and have been the focus of our discussions with students:

- Science seeks to explain and/or predict natural phenomena.
- Science offers testable ideas, and relies on evidence to judge ideas (confirmation/falsification).
- Scientific researchers form a community with certain standards (peer review; replication of tests)
- Science makes progress through time.
- Science tends to use mathematical or computational tools.

Further reading:

Stanford Encyclopedia of Philosophy - “Science and Pseudo-Science” https://plato.stanford.edu/entries/pseudo-science/

C. The Nature of Scientific Evidence

As noted above in the discussion of the demarcation problem, distinguishing what makes science special is not always easy. However, science does seem to importantly involve evidence, particularly in the form of experimentation and observation. Given this, it is important to consider what kinds of evidence are needed—and whether some kinds are better than others—and how much evidence we need to accept a given hypothesis or theory. In other words, even if we have “solved” the demarcation problem insofar as science – but not pseudoscience – relies on evidence, we have to approach the secondary questions regarding the nature of evidence. A full exploration of this topic isn’t possible, and the role of statistical reasoning in judging the support provided by evidence is beyond our current scope. But considering the nature of evidence and how it supports hypotheses is an important endeavor for understanding science.

Backing up a bit, the role played by evidence in science is a part of the larger question of how we come to know things in other disciplines and in everyday life. The study of knowledge – epistemology – is a major focus of philosophers. This begins with the notion of belief: a belief is the attitude one has when one takes something to be the case or to be true. It is a very general category that encompasses everything from casual opinions about things all the way to the kind of considered judgments that might underlie a firm statement of fact. Philosophers then ask: what is it that makes something one believes also something that one knows? In other words, what ingredients must be added to a belief for it to constitute knowledge? Traditionally, the first ingredient is that the belief must actually be true (a mistaken belief can’t be knowledge). Beyond this, however, more is needed, because a belief that happened to be true just by chance is not considered knowledge. The belief must also be something that we are justified in maintaining: perhaps because it is supported by reasons in the right way or gained via a reliable process. The exact way these latter ideas should be worked out is a subject for ongoing debate.

What constitutes good justification or support for a statement of fact or claim of knowledge? This can be another complicated topic, but we will focus on the idea that support relies on gathering supporting evidence of one kind or another. The most basic evidence is that which comes from direct observation. We trust that most of the time our senses give us reliable information about the world. With scientific evidence, such observations are often assisted by equipment (e.g. a microscope) that we have reason to believe is also reliable. Of course, we often don't make the observation ourselves, but rely on others. This brings up an interesting issue: when is it OK to rely on another person's information (or "testimony") when expressing a fact? Here, we tend to trust information if we believe the process by which it was obtained by others was sound. While the decision to trust testimony is not always easy, we generally seem justified in doing so if we understand in a rough way how it was acquired and have no particular reasons for doubting it. Here the institutions and norms that govern science play a major role. The trustworthiness of scientific sources of knowledge depends on training and certification of the expertise of individual scientists, but importantly also on self-checking mechanisms such as peer review of research publications and replication of experimental results. The operation of these and other social/institutional factors play a key role in assignment of authority to scientific testimony.

But when do we know if we have enough justification? This topic will be addressed in the next section. Here, though we can note that scientists are reluctant to declare that a hypothesis or theory is proved once and for all. The nature of scientific methodology is to continue to probe accepted ideas and wherever possible extend our understanding further. And, of course, the history of science shows many examples of theories that were superseded. This humility with which science proceeds should not, however, give rise to undue skepticism. Our best scientific theories have been extremely successful and have led to vast increases in our understanding while also forming the basis of reliable technologies. With regard to "proof", philosophers will usually say that it is not a notion that applies to claims about the world (it is a concept defined within logic and mathematics). But many established scientific facts and theories are so well supported by evidence that entertaining skepticism about them isn't reasonable. Saying something has been proven by science is often acceptable in everyday discourse, as long as one understands that this is not the strict proof of math or deductive logic: it is more akin to the legal notion of proof beyond a reasonable doubt.

The case of "proof" is one example of a contrast between the way we use words in everyday life and in specialized settings. Another prime example of this phenomenon worthy of noting concerns the word "theory." In everyday life, the word may indicate a not-well-supported statement about the world—some may use it to even mean just a hunch or guess about something. In scientific contexts, however, a theory is an idea (or set of ideas) that structure and explain a broad range of worldly phenomena. Importantly, a scientific theory is well supported by a variety of evidence – observations and/or experimental data about the various phenomena it describes and predicts.

In our class we explore the nature of scientific evidence through examples. These lead to discussions about a variety of factors, including:

- The importance of repeated testing and replication by other scientists. All else equal, more testing is better, including under a variety of conditions.
- The degrees of experimental control is important (including the reliability of experimental equipment). More control is better, all else equal, but there are often tradeoffs between control and other factors, such as sample size.
- The implications of using imperfect models (such as laboratory rats for medical research). Models that better approximate the target system of interest are preferred, but there is a tradeoff with control.
- The value of experimental tests vs. more indirect theoretical derivations. Sometimes we can use inference to extend our knowledge, but testing is still highly valued.
- The notion of multiple independent lines of evidence. A hypothesis or theory might be supported in relatively independent ways, and this is valuable, all else equal.
- The value of uncovering detailed causal mechanisms vs. finding what makes a difference to an outcome where the mechanism is not known. The former is often preferred but the latter approach can be powerful and sometimes easier to test.

In science, different considerations drive what evidence is potentially available and most relevant, so broad conclusions about the nature of evidence tend to be tentative. No single account of the amount/type of evidence needed is possible. Assessing the needs of each case is a kind of scientific expertise acquired in training and practice.

Further reading:

Internet Encyclopedia of Philosophy: “Epistemology” https://iep.utm.edu/epistemo/; “Evidence” https://iep.utm.edu/evidence/

D. Arguments and Scientific Inference

1. Introduction

A discussion of evidence and its importance to science prompts an interesting question: how exactly does evidence support a hypothesis? This question connects to the general topic of human reasoning and logical inference. In this section we take a structured approach to looking at how claims about the world are supported. This involves organizing our reasoning into arguments. An argument is a structured list of statements that present reasons (called premises) in support of a claim (the conclusion). When we gather evidence in an attempt to justify a claim, we can organize it into written statements expressing their contents—these are the premises of the argument—and evaluate the support that it provides for the claim—the conclusion of the argument. We will introduce two main kinds of reasoning (or rules of inference) that we might employ. These are deductive and inductive arguments. The latter is more common in scientific contexts, but it also poses a challenge: unlike in the case of a deductive argument, the premises of an inductive argument, even if true, do not guarantee the truth of the conclusion. The lesson also explores this issue and its implications (the "problem of induction").

Note that we make inferences throughout our daily lives. We reason from things we know or think are likely to be true to draw further conclusions. If the ground is wet outside our door, I infer that it has rained. From there, I might infer that my dog, who was left outside, will be muddy and in need of a bath. However, just as in the case of science, the claims we are seeking to support with arguments must have to do with something about the world or what is real, independent of our idiosyncratic perspective on things. In other words, the claim must have objective content. This is in contrast to other statements that depend on the subjective opinions or feelings of the person expressing them: such as when I insist to you that vanilla is the best flavor of ice cream.

2. Deductive Arguments

The first kind of reasoning is deductive, and uses deductive arguments. What is distinctive about a deductive argument is that the conclusion is supposed to be guaranteed to be true if the premises are true. Put another way, it is not possible for the conclusion to be false if the argument is structured correctly and has true premises. Here is an example.

Example (a):

Premise 1: All dogs are mammals.

Premise 2: Spot is a dog.

Conclusion: Spot is a mammal.

You might notice that there is a sense in which the content of the conclusion was implicitly contained in the premises (philosophers would say the argument is not ampliative), and this is why the truth of the conclusion follows from the truth of the premises. The form of this example, that proceeds from a more general claim (about dogs) to a claim about a specific instance (the dog named Spot), is one of several typical ways that deductive arguments proceed.

To evaluate a deductive argument that has been presented, consider that there are two ways that they can go wrong. First, the premises might not, in fact, guarantee the conclusion as they should. Second, the premises might not be true.

Example (b):

Premise 1: All dogs are mammals.

Premise 2: Spot is a mammal.

Conclusion: Spot is a dog.

Example (c):

Premise 1: All frogs are mammals.

Premise 2: Spot is a frog.

Conclusion: Spot is a mammal.

In example (b), the truth of the conclusion does not follow even if the premises are true. Spot, after all, could be a cat or horse (for the purposes of evaluating an argument we only use information presented to us—so it doesn't figure into things if there is by chance a dog in one's neighborhood named Spot). Finding alternatives to the stated conclusion that also follow from the given premises is an excellent way to show that an argument has failed. In the jargon sometimes used, we would say the argument is not valid, since in a valid deductive argument the truth of the premises ensures the truth of the conclusion. Turning to example (c): here the structure looks correct (the argument is valid), but Premise 1 is not true. In this case it is said that the argument is not sound. A deductive argument needs to have both the correct structure and true premises to be sound. And, of course, whether or not the premises are true will often be the focal point in evaluating an argument. If we have gathered reasons or evidence in support of a claim, and they do indeed justify its truth, attention may turn to how we can be sure that the supporting reasons are actually true.

Sometimes an argument seems like it should be valid, but some information seems to be lacking from the premises. Take the following example;

Example (d):

Premise 1: This distant star is over 10 billion years old.

Conclusion: The universe is over 10 billion years old.

Assuming the premise is true, the conclusion makes sense, but does not seem guaranteed: something is missing (what if the star somehow originated in a different universe somehow?). Compare this revised version of the argument.

Example (e):

Premise 1: This distant star is over 10 billion years old.

Premise 2: The universe must be older than any star within it.

Conclusion: The universe is over 10 billion years old.

Here we have made explicit a "hidden premise" that was in the background. Arguments need to make all the needed premises explicit for the structure to be valid. In the case of an argument seeking to support a scientific hypothesis, there are often background premises that might be left unsaid: for example, “our equipment was working properly”.

Deductive arguments do play a role in everyday and scientific reasoning, but we often want to use evidence to support conclusions that expand our knowledge beyond what may have been simply implicit in the premises. In such cases we will see, however, that guaranteeing a true conclusion is problematic.

3. Inductive Arguments

A second kind of reasoning uses inductive arguments. Unlike the case of a deductive argument, in an inductive argument the conclusion is not ensured by the truth of the premises, but is made likely to be (or probably) true.

Example (f):

Premise 1: Every swan that has been observed is white.

Conclusion: All swans are white.

Here, the conclusion moves from instances to a general claim (it is ampliative). Clearly, the conclusion could be wrong even if the premise is true (the next swan observed might be black – an example that has a real-life history: prior to the 17th century, Europeans had only seen white swans, and it was a surprise when they first observed black ones in Australia). The observations appeared to lend strong support for the conclusion, but they could not guarantee its truth.

As an aside, note that the falsification of a general claim *can* proceed deductively: from the premise that one black swan is observed, the conclusion that “not all swans are white” is guaranteed.

Scientific reasoning often relies on inductive reasoning to generalize (or to project into the future), often relying on statistical information. Their conclusions are therefore not guaranteed to be true in the same way as in deduction. It is more appropriate, then, to qualify the conclusion by claiming the premises make the conclusion “likely” or “probable”. This is the case even when the inductive conclusion appears quite strong. Consider this argument:

Example (g):

Premise 1: All living things examined require water for survival.

Conclusion: All living things probably require water for survival.

Despite the consistency of the observations, however, scientists will still be open to exceptions (some hypothesize that another substance might take on the role of water for an exotic form of life).

Note that, as in the case of deductive arguments, there are two ways an inductive argument can go wrong. Either the likelihood of the truth of the conclusion does not follow from the premises, or the premises themselves are false/unlikely to be true.

4. The Problem of Induction

The lack of certainty provided by inductive reasoning, however, can seem frustrating from our everyday perspective. Consider this classic example:

Example (h):

Premise 1: The sun has risen every day of my life.

Conclusion: The sun will rise tomorrow.

It may seem unreasonable to object that the conclusion is only probably true in a case like this one. After all, we rely on the consistency of patterns in the world to successfully navigate our lives. Perhaps we can add a premise to strengthen the conclusion? Consider this argument:

Example (i):

Premise 1: The sun has risen every day of our lives.

Premise 2: The past has been a good guide to prediction.

Conclusion: The sun will rise tomorrow.

But how do we justify this new premise 2? A skeptic will quickly point out that just because the past has been a good guide to prediction *so far* does not guarantee it will be so tomorrow. In effect we have used inductive reasoning in premise 2, so we have not strengthened our original inductive argument after all.

Relying on the uniformity of our past experience leads to a kind of circular reasoning, made explicit in this example:

Example (j);

Premise 1: Inductive reasoning has worked well.

Conclusion: Inductive reasoning will continue to work well.

We cannot guarantee our conclusions using inductive reasoning. However, as a pragmatic matter, it is very powerful and a solid basis for scientific conclusions. It is just that in dealing with the actual world, rather than the world of logic or mathematics, absolute guarantees are hard to come by.

Further reading:

Internet Encyclopedia of Philosophy: “Deductive and Inductive Arguments” https://iep.utm.edu/ded-ind/

Stanford Encyclopedia of Philosophy: “The Problem of Induction” https://plato.stanford.edu/entries/induction-problem/

E. Science and Morality

The demarcation problem was discussed in section B above, but here we look at a particular distinction: how does science relate to or apply to morality or ethics. In broaching this topic, we find that students bring with them a number of background ideas. They often think of morality descriptively: it is thought of as something embedded in a culture, or that is associated with religion. From these sources, we inherit certain common codes of conduct and prevailing attitudes about what is the right way to conduct oneself. It is helpful to introduce the idea that philosophers typically approach morality in a different way. They (prescriptively) seek general theories for how a person should act. Like in the case of science, moral theorizing doesn't assume the answers it seeks are relative to a particular person, group, or culture. There are many moral theories, and we will not review them here. But for background, it is helpful to briefly highlight two prominent approaches.

The first approach is consequentialism: a morally right act is one that leads to favorable consequences compared to other options, with what counts as favorable/unfavorable defined in various ways. Classic formulations focus on the goal of maximizing aggregate pleasure or happiness and minimizing suffering or unhappiness. With consequentialism, one must undertake a cost/benefit analysis of sorts to judge whether an act is morally right.

A competing family of theories says that we should not attempt to focus on consequences (which may be difficult to assess), but should define the proper moral rules or principles that our actions should follow. An example here would be the golden rule: in any given situation, act toward others as you would want them to act toward you. There may be a list of such principles that define a moral code. Note here that the golden rule, like other basic moral rules, is *universal* in nature. As is the case in the simplest forms of consequentialism, universal moral rules seek to offer the correct guide to everyone's moral conduct across all scenarios. They are not intended to be limited to a group or culture.

Science aims to describe, explain, and predict things in the world. Its hypotheses are about aspects of what is real, in other words, facts. What makes the relationship between science and morals difficult to assess is that it is controversial whether morality is also rooted in facts: many philosophers would deny this. And many of those who defend the idea of real moral facts would typically concede that, whatever these are, they are distinct from the natural facts pursued by science.

For these reasons, most would acknowledge that science cannot, by itself, answer moral questions. However, there is general agreement that scientific uncovering of natural facts might be of help in addressing them. This is perhaps easiest to see when consequentialism is the background moral framework (for instance: if we could agree that the right goal of morals is to maximize something like “well-being of sentient creatures” and if this is in principle measurable). But the broad conclusion, regardless of approach, is that science generally will not answer moral questions in the absence of some accompanying moral assumptions. Put in terms of a philosophical argument, we would say that scientific facts can be needed premises in argument for a moral conclusion, but they can’t be the only premises. These are distinctions that students generally have been capable of appreciating.

Further reading:

Internet Encyclopedia of Philosophy: “Ethics” https://iep.utm.edu/ethics/

Stanford Encyclopedia of Philosophy: “The Definition of Morality” https://plato.stanford.edu/entries/morality-definition/

Stanford Encyclopedia of Philosophy: “Consequentialism” https://plato.stanford.edu/entries/consequentialism/

F. Values and Objectivity in Science

1. Introduction

Objectivity as it relates to science is a complex topic. An initial distinction needs to be made between the kinds of claims that science deals in (e.g. hypotheses and theories), and the methods and processes for formulating and evaluating these claims and interpreting and applying the results. Scientific claims have objective content: they are meant to be claims about the world irrespective of any idiosyncratic human viewpoint. In other words, they concern facts. On the other hand, the notion of objectivity in the practice of science, in all of its phases, is fraught. After a brief clarification of the distinction between facts and opinions, we will turn to discussion of how human biases and values intersect with scientific practice.

2. Facts vs. Opinions

Perhaps someone has responded to something you said with this question: “Is that a fact or is it just your opinion?” We often see “fact” and “opinion” opposed to each other in this way. If we take a closer look, however, we can see that the two options do not neatly divide the conceptual terrain we care about when we ask this question.

Philosophers consider facts to be things that actually exist; they are aspects of what is real. So, to be clear, the question is asking about whether the content of something you have said (perhaps relating a belief you hold) states or expresses a fact. Asking if I have stated a fact is asking whether what I have said corresponds to something that is the case, or is real. If it does, then, I have made a true statement about the world (according to a simple conception of truth - more on this below).

Note that another (more directly opposing) contrast to my stating a fact would be if I had said something *false* about the world. These are the two options if the statement is about something that describes how the world might be. It either expresses a fact, and is true, or, on the other hand, it expresses something that is not a fact and is false. In either case, what one is expressing must be a claim about the world that doesn't depend on my idiosyncratic perspective as a human being. This contrast can be described in terms of statements having "objective" versus "subjective" contents. Both true and false statements about the world have objective content: their truth or falsity doesn't depend on my particular subjective point-of-view.

With this as backdrop, we can now see that offering an "opinion" doesn't neatly contrast with giving a statement of fact (an opinion is not defined as the expression of a falsehood). In everyday usage there appear to be *two different things* one might mean by "opinion" when asking "Is that a fact or is it just your opinion?" First, it might refer to a statement or claim that purports to be true of the world, but might well be false. Here, the label of "opinion" is expressing doubt about the statement's truth. But there is a second possibility: perhaps it is an "opinion" because it isn't the kind of thing that could be either true *or* false in a way that is independent of my human perspective (it expresses something subjective).

Let's look at an example of the first usage. I stub my big toe badly and declare to my friend: "My toe is broken!" She says: "You don't know it's broken—that's just your opinion." Now, whether my toe is broken is the kind of thing that is either true or false about an aspect of the world. My friend is suggesting, however, that my statement might be wrong: she is thinking I do not have good reasons in support of its truth (after all, I am no expert). Now assume that I then visit the emergency room. The doctor examines my toe and takes X-rays, and subsequently concludes my toe is indeed broken. Having relayed this news to my friend, she now considers it a fact that "my toe is broken". The statement has been appropriately justified or supported (by an experienced doctor using appropriate tools and methods).

In contrast, assume I say: "Vanilla ice cream is much better than chocolate." Now if I am accused of merely offering an opinion, this is for a different reason. The statement has no truth or falsity apart from my perspective: it is subjective. As a simple test, the difference between the two uses of opinion can be judged by whether we can imagine it to be true or not independent of the beliefs and feelings of particular human beings. In the broken toe case, my particular feelings played a role in my making the initial assertion (my toe is broken!), but my beliefs and feelings played no role in determining the statement's truth or falsity. In the ice cream case, it is my personal feeling or preference that is being directly expressed.

Some cases may be harder to classify. I say: “Big dogs are much scarier than small dogs.” This may be seen as an opinion of the second sort: it is subjective. However, it is likely a bona fide fact that big dogs are capable of causing greater harm to a human. To the extent this is the idea I am imperfectly conveying, perhaps one might say I am making a well-supported claim about the world (just not very clearly).

3. A Note on Truth

Science is thus concerned with statements expressing facts, and whether their truth is adequately supported. But philosophers sometimes worry about whether the distinction made above between claims with objective content (that might be true or false about the world) vs. those with subjective content (that rely on a particular point of view) is oversimplified. Can a statement (or proposition) be considered to be true or false independently of the existence of humans? One issue to consider is our human fallibility in determining whether our statements correspond to facts. And a further problem perhaps lies in the assumption that we can even conceive of worldly facts without reference to human thoughts and language. There are, in fact, many philosophical theories of truth in addition to the correspondence account implied in this discussion (and some versions of the correspondence theory allow that facts may be in some way dependent on human minds). For many everyday examples and, importantly, for science, however, it seems sufficient to work with the idea that our statements are true if they convey what is the case in the world independent of human minds.

4. Human Values

How can science deal in the currency of facts if it is an activity conducted by humans with various biases and preferences? Here we draw a distinction between the objective content of the hypotheses, theories and conclusions of science on the one hand, and the objectivity of the methods and processes employed. Can the practice of science be free of the influence of human values or biases? Is this the right goal, even if difficult or impossible to achieve?

For brevity's sale, we will use the term "values" as a broad category that encompasses an individual's moral/ethical stances, their cultural/religious/social/political viewpoints, and personal preferences and biases. These are the kinds of values that would seem potentially problematic. It is worth briefly mentioning that there is a different set of so-called epistemic or cognitive values that are traditionally acknowledged as being an accepted part of science. These are preferences that scientists might have regarding theory choice, when two or more theories are compatible with the evidence that has been gathered. Examples of these are simplicity, scope of application, and coherence with other accepted theories. We are focused here on the role of the more controversial, "non-cognitive" values listed earlier.

In exploring this topic, we do not focus on cases of scientific fraud. Every so often, someone in search of professional success fakes data or otherwise subverts the scientific process in an obvious way. Eventually the institutional structures of science (peer review, examination by other researchers and labs working in the same field) tend to uncover these problems. It is the more subtle and potentially unacknowledged role of values on the scientific process that is of greater interest.

5. The stages where values may enter science

To understand the actual and potential impact of values on science, it is necessary to consider the different stages of the scientific process where values might enter. Here is one way to break things down (with the important caveat that this is one of many ways of doing so: the actual process of science is more complex and includes variations and iterations):

- Choosing a Scientific Research Topic
- Allocating Funding to Research
- Designing Experiments/Choosing Research Methods
- Interpretation of Experimental Findings/Deciding to Accept or Reject Hypotheses
- Deciding Whether/How to Publish/Publicize Results
- Choosing Whether/How to Apply Research Results to Practical Problems.

Beginning with the first stage: personal values undoubtedly play an important role in the pursuit of research topics. The initial decision to enter a scientific field, and then to focus on certain topics or questions are based on a person's interests and goals. These value influences seem at a minimum to be benign. In the cases of scientists wanting to uncover new results that benefit other human beings or the environment, their aims are laudable. But even in cases where the goal is something like career advancement or higher salary, there is nothing necessarily harmful per se with letting values play a role at this stage. Of course, there may be a big picture societal interest in making sure some topics are pursued or at least not neglected. But this will naturally concern values at higher collective or institutional levels.

On the other hand, such values, which enter into our second stage (allocating funding to research), can be problematic. In theory government can play a positive role in allocated resources to address societal needs, appropriately guided by democratic values. But sometimes these values are questionable in practice, and the large role played by corporate and other private interests in funding research is also of great concern. Cigarette companies, purveyors of unhealthy foods, and defense contractors have all funded research that has the potential to perniciously influence the practice of science. From the point of view of the individual scientists accepting such funding (in a setting like a university), the potential for the funder's values to bias the results (even if unconsciously) is real and potentially harmful.

The third stage is one that perhaps gets less attention. But scientists, having chosen a research goal, have many subsidiary decisions to make to spend their funding and implement their research program. Decisions such as who to hire in the laboratory and what research subjects and equipment to use might be influenced by personal biases or even just a desire to take the easiest path toward finding a plausible result to advance one's career (using laboratory rats for medical research; using undergraduates in psychological research).

The fourth stage is the one most people probably have in mind when thinking about the need for scientists to be value-free in pursuing their work. The interpretation of results and decision to accept or reject hypotheses shouldn't be influenced by values if we want our research to uncover facts. While it is true that there is great risk here, it is worth noting that, even at this stage, it might be possible that values can enter in an arguably benign way. For instance, consider a case where a scientist is interpreting results where there is an asymmetric risk involved in how the results will be applied to a related policy question. For instance, assume a chemical being discharged into the water supply of a community might cause cancer; preventing its discharge would be only moderately costly. To test the question, a large number of lab mice are fed the chemical. Later, some proportion of them do get cancer and the question is whether or not to reject the hypothesis that the incidence is within normal bounds. Further, assume that when examining tissue samples drawn from the mice through a microscope for signs of malignant tumors, some of the samples are ambiguous (hard to tell if cancer is present or not). Would it be appropriate for the researchers to categorize these cases with the clearer cases of malignant tumors, since they believe the consequences of underestimating cancer risk are worse than those of overestimating it? Also, if a statistical significance test is applied to judge the outcome of the research, should this be weighted in favor of risking a false positive signal (cancer is erroneously thought to be present in greater than expected amounts) rather than a false negative, when compared to the standards used in experiments where no such asymmetric risk exists?

In thinking about a case like this, many would still lean toward thinking that answering yes to these questions would be inappropriate intrusions of values into the scientific process. The scientist should just conduct the research as neutrally as possible. But perhaps such an intrusion of values into the interpretation and hypothesis evaluation steps would be acceptable if these modifications to the process were fully disclosed by the researchers and their likely effects on the outcome transparently documented. We will return to this issue below in the concluding section.

The fifth stage above was deciding whether or how to publish or publicize results. The problems that might enter at this stage have become more widely discussed recently. It is a concern if biases exist favoring the submission and publishing of some kinds of results compared to others. For example, in many fields there seems to be a systematic bias in favor of reporting and publishing positive results (a hypothesis is tested and confirmed) vs. negative results (an experiment failed to reveal the expected result). The bias of scientists and scientific journals in favor of interesting positive results may have a negative side effect, as research productivity would be aided if all negative results were likewise publicized (so that dead ends are not pursued unknowingly by other researchers).

The role of values in the final stage listed above is not generally controversial. How we use scientific knowledge in our lives will depend on our goals. In a democratic society, we hope and expect that important policy decisions are informed appropriately by science.

6. Conclusions

Surveying the ways values can enter the scientific process leads us in the direction of two conclusions. First, while science does indeed deal with facts, insisting on the process being completely value-free is unrealistic, and sometimes inappropriate, since values can enter in a neutral or positive way at some stages. The second key point is that when the influence of values is indeed potentially problematic, the greatest risk to the process is *unacknowledged and/or unreported* bias. The social and institutional structures of science (university training, collaboration/friendly competition between researchers, peer review) work their best when all the possible value judgements and biases are disclosed.

As students, we are sometimes taught a simplified picture of the scientific process, where practitioners and their methods are depicted as completely objective. The real situation is more complex, but understanding it also offers the potential for greater confidence. Science has strong institutional structures that allow it to mostly deliver accurate and helpful descriptions, explanations and forecasts. But all of this occurs through the work of human beings who hold to various values, just like the rest of us.

Further reading:

John Corvino: "The Fact/Opinion Distinction" https://www.philosophersmag.com/essays/26-the-fact-opinion-distinction

Stanford Encyclopedia of Philosophy: "Scientific Objectivity" https://plato.stanford.edu/entries/scientific-objectivity/