Part I

Background and Basic Concepts

1 Some problems of induction

1.1 Introduction

To the lay reader, much of what is written by scientists can seem barely comprehensible. Even to someone who has had some science courses in school, a sentence like "The M2 proton channel ... has a 40-residue region that interacts with membranes consisting of a transmembrane helix (which mediates tetramerization, drug-binding, and channel activity), followed by a basic amphiphilic helix important for budding of the virus from cellular membranes and release (scission)" will seem as though it has been written in a language not quite identical to English (Fiorin, Carnevale, & DeGrado, 2010). As a result, the non-specialist may find much of what is accomplished by scientists mysterious and esoteric.

Philosophers of science have sometimes attempted to "humanize" scientific work by portraying scientific methods as extensions of our ordinary ways of knowing things. It is true that scientists use technologies both material and mathematical to make observations and draw conclusions that we could never achieve otherwise. Nonetheless, they observe, conjecture, infer, and decide just as we all do, if perhaps in a more systematic and sophisticated way. Although a certain kind of training is needed to understand some of the language scientists use to report their findings, those findings are not the result of magic or trickery, but of an especially refined and disciplined application of the cognitive resources we enjoy as a species.

But if scientists use and extend the cognitive abilities of ordinary knowers, they also inherit the philosophical problems of ordinary knowers. One of the thorniest and most discussed of these – the problem of induction – is the topic of this chapter. We will see how this general problem about knowledge arises in both non-scientific and scientific contexts. An example from the history of optics will show that scientific experimentation and careful use of scientific methods can improve our ability to give well-reasoned answers to questions, but do not suffice to solve the problem of induction itself.

1.2 The problem of induction about doughnuts

The problem of induction is the problem of how we learn from experience. Consider how it arises in the life of a somewhat ordinary person, whom we shall call Zig. Zig likes doughnuts, let's suppose, and one day decides to try a new doughnut shop. This doughnut shop has introduced a new horseradish-flavored doughnut, something that Zig has never tried before. Zig expects, however, to like the horseradish-flavored doughnut. After all, he has enjoyed nearly all of the doughnuts he has eaten in his life (he has eaten many!). In fact, Zig believes there to be a general Law of Doughnut Delectability:

Law 1 (*LDD***)** *All doughnuts are delicious.*

Since horseradish doughnuts are doughnuts, it follows that horseradish doughnuts must also be delicious.

Zig's reasoning here exhibits a logical trait of considerable significance. He is using an argument that is *deductively valid*, meaning that *it is impossible that the premises are true and the conclusion false*. This can be seen clearly enough in the case of Zig's argument. Suppose the *LDD* is true, and that it is true (as it seems it must be) that all horseradish doughnuts are doughnuts. Could there possibly be a non-delicious horseradish doughnut? Suppose there were. If we insist on the law that all doughnuts are delicious (first premise), then this non-delicious horseradish doughnut must not really be a doughnut after all. On the other hand, if we insist that all horseradish doughnuts really are doughnuts (second premise), then our non-delicious horseradish doughnut stands in contradiction to the law that all doughnuts are delicious. We simply cannot hang on to the truth of our premises while denying the conclusion of this argument.

Deductive arguments are sometimes said to be 'truth-preserving' because a deductively valid argument is guaranteed not to lead you from true premises to a false conclusion. This makes them powerful as logical tools of persuasion; as long as the person you are trying to convince agrees with you that your premises are true, she cannot resist the conclusion of a valid argument from those premises without contradicting herself. A deductively valid argument with true premises is said to be *sound*.

To his very great disappointment, Zig finds the horseradish doughnut revolting. He has just encountered a limitation of deductively valid arguments. Although the truth of their premises guarantees the truth of their conclusions, there is nothing about validity itself that guarantees the truth of the premises. Zig proceeds to make a new deductively valid inference: All horseradish doughnuts are doughnuts. Some horseradish doughnuts are not delicious (to say the least!). Therefore, some doughnuts are not delicious. With this argument, Zig has refuted the *LDD*.[1]

Now Zig begins to wonder: was that one horrendous-tasting horseradish doughnut a fluke? The next day he returns to the doughnut shop and orders another one. The second horseradish doughnut turns out to be equally vile. After ten more visits to the same shop, each time ordering a horseradish doughnut that he finds quite disgusting, Zig begins to visit other doughnut shops selling horseradish doughnuts. Zig proceeds to sample 12 horseradish doughnuts from each of 20 different doughnut shops and finds all 240 of them to be repellent.

Surely Zig is overdoing it! He has, you would think, more than enough evidence to justify an expectation that *any* horseradish doughnut he tries will not taste good to him. But suppose he remains skeptical of this claim. He freely admits the unpleasant taste of all horseradish doughnuts he has eaten in the past, but does not think that this has any bearing on what he should expect of those he has not yet tasted. What logical argument could we offer to persuade him?

The problem of induction amounts to just this skeptical problem: it seems that no amount of unpleasant experience with horrid horseradish doughnuts would suffice to justify any belief regarding the deliciousness or putridity of one that is yet to be tasted.

Consider the premises we have at our disposal for reasoning with Zig. "Look, Zig," we might say. "You have eaten 240 horseradish doughnuts from a variety of doughnut shops in a variety of locations. You agree that each of those doughnuts was deeply unpleasant. Does it not follow that the next horseradish doughnut will be similarly awful?" Zig would be correct to reply "No, it does not follow. The premises you offer are consistent both with the

[1] You may have noticed that Zig could have saved the *LDD* if he had been willing to consider giving up the claim that all horseradish doughnuts are doughnuts. Maybe the terrible-tasting thing that he ate was not really a doughnut after all? We will return to the possibility of such evasions in Chapter 3.

next one being tasty and with it being terrible." In other words, the suggested inference to the conclusion that all horseradish doughnuts are not delicious (or even that just the next one lacks gustatory appeal) is not deductively valid. Our inference regarding the next horseradish doughnut is an inductive, not a deductive argument, and as presented it does not exhibit validity.

One might try to make it into a deductive argument by filling in some additional premises. Consider, for example, the following argument: All horseradish doughnuts in a large and varied sample are not delicious. All horseradish doughnuts that have not been tasted are similar to those in the sample with respect to their taste. Therefore, all horseradish doughnuts are not delicious.

Now we have a valid deductive argument, but we also have a *new premise*, and it is one that cannot be considered to report any observation that Zig has made. 'All horseradish doughnuts that have not been tasted are similar to those in the sample with respect to their taste' asserts a generalization about horseradish doughnuts, and Zig would rightly ask us to tell him what reasoning would justify his accepting this new premise. It is no more a deductive consequence of the uniformly bad taste of the horseradish doughnuts Zig has already eaten than is our original conclusion.

Perhaps this new premise can be supported by some separate argument that we think Zig should find compelling. One could, for example, regard it as a special case of a general Principle of Uniformity of Doughnut-Properties (PUDP): All doughnuts of a given kind resemble one another with regard to their taste. And this general principle, we could argue, is supported by past doughnut-eating experiences, not just of Zig, but of the general doughnut-eating public, which has found in the past that all raspberry-filled doughnuts tasted similar, all chocolate cake doughnuts tasted similar, all sour cream doughnuts tasted similar, etc.

Clearly, however, this justification for the PUDP is itself an inductive argument. Having declined to accept the conclusion of our original induction, Zig is under no logical compulsion to accept the conclusion of this argument either. We could try to patch up the logical gap with some other general principle such as the Uniformity of Food Properties, but then *this* principle will need to be supported by an inductive argument from uniformity in past food experiences.

At each stage in our effort to close the logical gap between premises and conclusion in an inductive argument, we will find ourselves invoking a principle that in turn has to be justified inductively. Now either this series of justifications ends or it does not. If it ends, then it must end with us appealing to some general principle for which no further inductive justification is offered. This most general such principle is sometimes simply called the Principle of the Uniformity of Nature (PUN). Just how to formulate the PUN so that it is both strong enough to perform its logical task and plausibly true is difficult to say.[2] But we do not need to formulate it to see the problem: It is a general principle that cannot simply be true by definition (like 'carnivores eat meat') or be derived deductively from other premises that are true by definition. If it does not lend itself to such a deductive proof, and we have no inductive argument to support it, then it must simply be unjustified by any argument whatsoever, and so then are any of its consequences.

If, on the other hand, the series of justifications does not end, then we have an infinite regress of inductive arguments, so that the justification for our original conclusion can never be completed. Either way, the conclusion of any inductive argument (for there is nothing special about our doughnut example) turns out to lack any justification.

Note that in this illustration of the problem of induction we did not equip Zig with any specialized methods or concepts of the sort that scientists bring to bear on their research problems. The problem of induction appears not to be a problem of science so much as it is a problem of knowledge in general.

1.3 Reasoning about light: Isaac Newton's induction

Here, however, let us set aside the problems posed by our imaginary doughnut investigator's skepticism, and consider an actual bit of historically significant scientific reasoning. What we will see is that although the investigator – in this case, Isaac Newton – gave no explicit attention to the problem of induction in this general form that we have just considered, he devoted considerable attention to the question of just what he could justifiably infer from his observations, and the assumptions that would be

[2] The difficulties can be seen already emerging in the PUDP: How strongly do the doughnuts in a given kind resemble one another? If the principle asserts that they are exactly alike, then the PUDP seems to be false; if they are only somewhat similar in a very broad sense then the principle will be too weak to support the conclusions that are meant to be drawn from it.

required to justify those inferences. We might say that, although he did not concern himself with *the* problem of induction, he did concern himself with a number of specific problems of induction, i.e., problems connected with particular ways in which he might go wrong in drawing his conclusions.

> In a very dark Chamber, at a round Hole, about one third Part of an Inch broad, made in the Shut of a Window, I placed a Glass Prism, whereby the Beam of the Sun's Light, which came in at that Hole, might be refracted upwards toward the opposite Wall of the Chamber, and there form a colour'd Image of the Sun.

Thus begins Isaac Newton's "Proof by Experiments" of his second proposition in his 1704 work *Opticks*. The proposition in question states that "The Light of the Sun consists of Rays differently Refrangible" (Newton, 1979, 26). In other words, sunlight can be decomposed into rays of light that differ from one another in the angle at which they are refracted through a given medium.[3]

Newton proceeds to describe how he arranged his prism so that its axis was perpendicular to the ray of light incident upon it, and so that the angles of refraction at the incident face of the prism and at the face from which the rays exited the prism were equal (see Figure 1.1). He describes the image made by the refracted light on a sheet of white paper as "Oblong and not Oval, but terminated with two Rectilinear and Parallel Sides, and two Semicircular Ends" (Newton, 1979, 28), and proceeds to report in detail on its dimensions. He reports that when the refracting angle of the prism was reduced from its initial value of 64 degrees, the length of the image was reduced, but its breadth was unchanged. If the prism was turned so as to make the rays exiting the prism emerge at a more oblique angle, the image was lengthened. He notes some irregularities in the glass of his first prism, and thus tries the same experiment with different prisms "with the same Success." "And because it is easy to commit a Mistake in placing the Prism in its due Posture, I repeated

[3] Newton had first reported on his experiments with prisms in his inaugural lectures as the newly appointed Lucasian Professor of Mathematics at Cambridge in 1669 (Newton, 1984). His first published paper on the topic appeared in the *Philosophical Transactions of the Royal Society* in February 1671. There he describes the experiments as having been performed five years earlier "in the beginning of the year 1666," many years before writing the *Opticks* (Newton, 1978). The correct historical sequence of Newton's optical experiments and the trustworthiness of his narrative are matters of historical dispute (Guerlac, 1983; Newton, 1984).

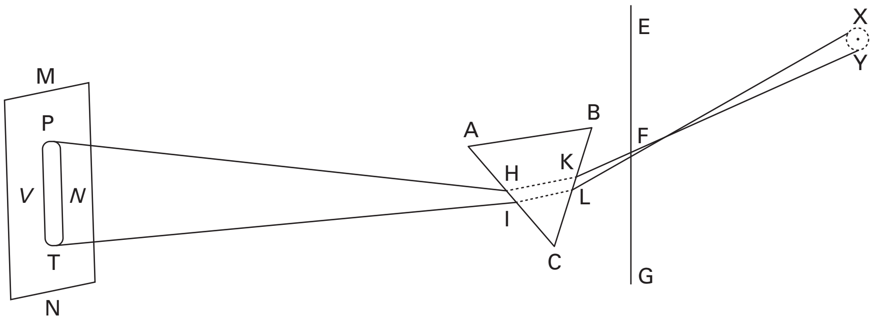

Figure 1.1 Newton's first experimental arrangement. *F* is the hole through which light from the sun (*XY*) passes, on its way to the prism (*ABC*). On a sheet of paper (*MN*) the light casts an image (*PT*) of the sun. (From Newton 1979, 27; by permission of Dover Publications.)

the Experiment four or five Times, and always found the Length of the Image" to correspond to that which he first reported (Newton, 1979, 30).

Newton then draws an intermediate conclusion from his measurements: The rays exiting the prism diverged from one another in such a manner as to constitute the length of the image, forming an angle "of more than two degrees and an half" (Newton, 1979, 31). That result Newton shows to be incompatible with "the Laws of Opticks vulgarly received," according to which the angles at which the rays exit the prism should diverge at an angle corresponding to the angle formed by the rays arriving at the hole from the extreme points of the sun's diameter: about half a degree. This would result in an image on the sheet of paper with a length equal to its breadth. What has gone wrong?

Newton points us to the answer by noting that the expectation of a round rather than elongated image would be correct were the rays that constitute the image "alike refrangible." "And therefore seeing by Experience it is found that the Image is not round, but about five times longer than broad, the Rays which going to the upper end . . . of the Image suffer the greatest Refraction, must be more refrangible than those which to go the lower end" (Newton, 1979, 32).

There follows in the *Opticks* Newton's description of a series of experiments involving various arrangements of prisms through which Newton explores the phenomenon just described. Here I will just describe one such permutation that seems particularly relevant for understanding Newton's defense of his proposition that the sun's light consists of rays that differ in their

Cambridge University Press
978-0-521-11249-9 - An Introduction to the Philosophy of Science
Kent W. Staley
Excerpt
More information

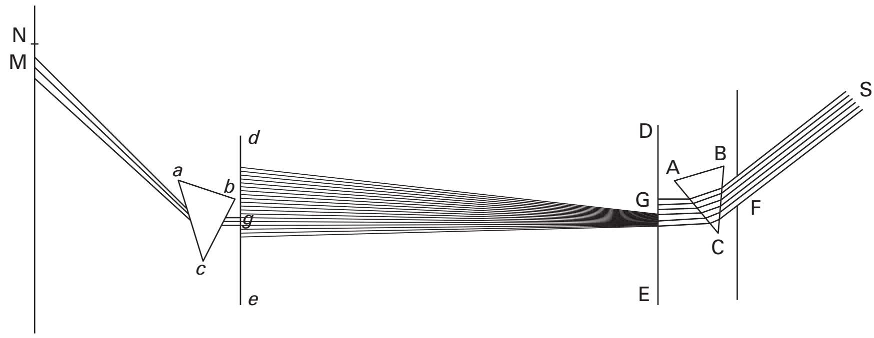

Figure 1.2 Newton's two-prism experimental arrangement. *S* is the Sun, light from which passes through an aperture *F* to prism 1 (*ABC*). A second aperture (*G*) can be used to select a portion of the refracted image of *S* to be sent through a third aperture (*g*) to prism 2 (*abc*). The refracted light from *abc* forms an image on the wall *MN*. (From Newton, 1979, 47; by permission of Dover Publications.)

refrangibility. In this arrangement (see Figure 1.2), Newton allowed sunlight to enter through "a much broader hole" before passing into a first prism, the orientation of which about its axis could be adjusted. The refracted light is then directed at a board with a hole about one-third of an inch in diameter. The beam of light thus formed is then projected onto a second board "and there paint[s] such an oblong coloured Image of the Sun" as had been found in the previously described experiment. This second board also has a small hole in it, and light that passes through that hole is then sent into a second prism, with its axis parallel to that of the first prism, which is held at a fixed position.

By varying the orientation of the first prism, Newton effectively sends different parts of the spectrum created by the first prism into the second prism. In this way he is able to more stringently focus on the question of whether these rays that emerge at different angles from the prism really do, *considered separately from one another*, refract at different angles. To test whether this is so, he observes the location on the opposite wall of the room at which the resultant ray of light arrives, having passed through both prisms. He found that indeed light from the higher part of the spectrum projected onto the second board refracted through the second prism to a higher location than light from the lower part of the spectrum, while light from intermediate parts of the spectrum arrived at intermediate locations on the wall. Newton thus is able to show that, although each ray arrives at the second prism at the same