Supplementary Section 7S.8

Logic and Science

Philosophy of science became a specialized subfield of philosophy concomitant with the development of modern, formal logic. Especially in the early twentieth century, logic played an important role in attempting to work out the nature of explanation and the related notion of confirmation. In this section, we will look at the relation of logical deduction to these topics and how logic may help us to understand the notion of a law. We will also see how the notion of a contradiction plays a central role in constructing and revising our scientific theories.

SCIENTIFIC EXPLANATION AND THE D-N MODEL

Explanation often begins with why-questions. Why do you study logic? Why did the United States enter World War I? Why does Earth revolve around its axis? Answers to such questions are often called explanations. Such questions solicit descriptions of the world that, in some sense, explain the event in question. Science is, at least in part, a collection of explanations, organizations of our best descriptions of the world.

There are many competing theories of explanation, a topic of lively debate in the philosophy of science. Among the earliest developers of formal theories of scientific explanation was Carl Hempel. In a series of papers including the classic "Studies in the Logic of Explanation," written with Paul Oppenheim, Hempel developed what he calls the deductive-nomological, or D-N, model of scientific explanation.

'Nomological' means law-like. According to the D-N model, an explanation of an event or phenomenon is a logical inference that uses general laws and initial conditions as premises and the phenomenon to be explained, the explanandum, as a conclusion. D-N explanations are deductive in that the explanandum is derived from the laws and initial conditions, as in the general form at 7S.8.1.

7S.8.1	$L_1, L_2, L_3, \ldots L_n$	The relevant laws . . .
	$I_1, I_2, I_3, \ldots I_n$	. . . and the relevant initial conditions . . .
	————————	. . . logically entail . . .
	E	. . . the explanandum.

Consider the explanation of why a batted baseball takes a particular route to left field. Among the general laws involved are those governing the transfer of momentum from the bat to the ball and the force of gravity holding between Earth and the ball. Among the particular conditions include the particular angle at which the bat hits the ball, the masses of the bat and ball, the wind speed, and the air pressure. Given the velocities and masses and angles of impact, one can infer the trajectory of the ball. Such an inference is, according to Hempel and other proponents of the D-N model, an explanation of the ball's flight.

Or consider why a person drinks a glass of water. In this case, the initial conditions might include the drinker's thirst and the presence of water and a cup. The general laws would include physical and psychological generalizations about behavior: if I am thirsty, and there is water available, I drink some; I am thirsty, and there is water available; so I drink.

Notice that different kinds of laws may be used in D-N explanations. In the batted-ball case, the laws to which we appeal are fundamental physical laws of motion. In the water-drinking case, the laws to which we appeal are psychological generalizations. The D-N model can provide a variety of kinds of explanations.

In ordinary cases, the laws involved in D-N explanations of particular events are general. They are sometimes called covering laws, since they apply to a wide range of cases and subsume particular events under them. The laws say that every time certain circumstances are realized, certain specific phenomena will occur. If we want to know why this sea creature is a mammal, we can appeal to a general law that all whales are mammals and the specific fact that this sea creature is a whale.

Both the general laws and the specific conditions subsumed under the general laws are required for a D-N explanation. Without specific conditions, we are left without application of the laws in any particular case. Without general laws, we lack an explanation of a particular event or phenomenon.

For example, we might try to infer that a baseball flew out to left field solely from the angle at which it was hit and the momentum of the pitched ball and swung bat. But the explanandum follows neither logically nor conceptually without the invocation of laws that govern or describe interactions of balls and bats.

Moreover, the mere presence of a deductive inference that yields an explanandum is insufficient to provide a D-N explanation. We could infer, logically, that the ball flew to left field from the premise that the ball was struck and flew to left field, as at 7S.8.2.

7S.8.2 The ball was struck and flew to left field.
Therefore, the ball flew to left field.

The inference at 7S.8.2 is perfectly deductively valid. It requires only the use of the propositional logical rule of simplification. Such an inference is not explanatory, though, but disappointingly circular, lacking the relevant general laws.

Such disappointing or vacuous explanations may be disguised by fancy language. Consider the mocking claim, found in *Le Malade Imaginaire* by the seventeenthcentury French playwright Molière, that opium puts one to sleep because it has the dormitive virtue. Molière, in invoking this faux explanation, was showing how people sometimes pretend to explain by constructing obscure terms. Saying that opium puts one to sleep because it has the dormitive virtue is to say that it puts one to sleep because it puts one to sleep, another disappointingly empty explanation. A better explanation of why opium puts one to sleep would appeal to general laws about how opium interacts chemically with our body. Such an explanation would fit the D-N model by offering general principles of chemistry and neurology and particular facts about opium and brains and bodies.

So the D-N model can be used to explain specific events or phenomena by using general laws. We can also use the D-N model to explain the so-called lower-level (or more-particular) laws by higher-level (or more-general) laws. For example, consider the three chemical laws at 7S.8.3.

7S.8.3	Boyle's law	$P_1V_1 = P_2V_2$
	Charles's law	$V_1/T_1 = V_2/T_2$
	Ideal gas law	$PV = kT$

Boyle's law and Charles's law are more lower level than the ideal gas law, more specific and more narrow. For Boyle's law, we assume a constant temperature while varying pressure and volume. For Charles's law, we assume a constant pressure and vary temperature and volume. The ideal gas law combines the results of the other two and applies more generally, in cases of varying temperature or varying pressure. And we can explain the lower-level laws in terms of the higher-level law, by holding either pressure or temperature constant and deriving the lower-level laws from the ideal gas law.

Hempel uses the example of how gravitational law explains Galileo's law regarding free-falling bodies. Galileo took the acceleration of a free-falling body to be a constant. Newton's law of gravitation shows that Galileo's law is false. Since the acceleration of a body due to its gravitational attraction to another body varies with the distance between the two bodies, as a falling object approaches Earth, its acceleration increases. But because of the size of Earth, Galileo's formulation approximates the better Newtonian formulation and can be used in many practical cases. The Newtonian law of gravitation can explain why Galileo's formulation works near the surface of Earth.

Thus, having a D-N explanation can be useful in explaining both particular phenomena and particular laws.

Several problems arise for using the D-N model to produce the best explanations. For many purposes, we want to discover the highest-level, the most general, laws available. In particular explanations, though, we generally do not need a high-level covering law and can use a minimal covering law instead. Newton's gravitational law, for example, applies to interactions among all massive bodies. But to explain the falling of, say, my keys on the floor, we do not need the broad laws of Newtonian gravitation or even the more correct, even higher level, relativistic formulations of gravitational theory. We can, instead, appeal to less-general laws that apply only to, say, the gravitational pull of Earth on these keys rather than the attraction of any two objects. Such a minimal covering law would provide a D-N explanation. The laws would refer only to the attraction of particular keys and the planet Earth. The initial conditions would refer to the relevant keys. The explanandum could be logically inferred. But the explanation may appear to be unsatisfying because the minimal covering is not general enough. Again, merely having a D-N form does not suffice to make an inference into a satisfying explanation.

A similar challenge for defenders of the D-N model of explanation concerns the direction of explanation. To take a classic example, consider a flagpole and its shadow, given the height of the sun. From the general laws of light, and the height of the flagpole, we can infer (and thus, explain) the length of the shadow. But we could equally well infer the height of the flagpole from the length of the shadow. While it seems reasonable to explain the shadow on the basis of the height of the pole, it seems odd to explain the height of the flagpole on the basis of the shadow. (We might explain our knowledge of the height of the flagpole by our knowledge of the length of its shadow, but that's a separate case.) Some conditions on the direction of explanation are required for the D-N model.

Other problems for the D-N model of explanation concern determining what the laws are. So far in the section, we have mainly been thinking about some uncontroversial cases of laws, like Newtonian gravitation or the ideal gas law. But we also saw that other kinds of generalizations play the roles of laws in D-N explanations, as when we explain why someone drinks a glass of water. And not all generalizations are laws, even if they look like laws in being general claims.

Consider the explanation of why a particular student in my classroom is under sixty years old. We could appeal to the fact that the student is a person in my classroom. And we could invoke the generalization that everyone in the classroom is under sixty. We could then infer, logically, that the student is under sixty, as at 7S.8.4.

7S.8.4 All persons in my classroom are under sixty.
Rey is a person in my classroom.
So Rey is under sixty.

7S.8.4 has the structure of a D-N explanation, and its first premise is a general claim that plays the role of the law in the explanation. But that general statement is not, in other ways, law-like. If a person over sixty came in to the room, for example, the purported law would no longer hold. In the baseball case, whether the bat hits the ball or not, the laws of transfer of momentum remain. In the classroom case, the general statement about the ages of people in the room holds only when particular people are in the room. One way to put this difference is that real laws support counterfactual instances.

The problem of determining which statements are laws and which are not is deep and difficult. One difference is that laws remain true no matter what particular events take place. But it is difficult to determine the truth of counterfactual claims just because they are counterfactual. Moreover, a solution to the problem seems essential to understanding D-N explanations because of the roles played by laws in such explanations. But philosophers not only disagree about how to determine the laws, they do not all agree about what a law is. Some philosophers think of laws as fundamental aspects of the world that govern, in some sense, events and interactions. Others think of the laws as mere summaries of interactions, having no reality of their own.

Though we will not pursue a solution to these problems here, two points are worth making. First, the problem of determining the laws (if there are any) is closely related to the problem of induction discussed in section 7.1 of IFLPA. Second, and relatedly, there is no syntactic test for whether a statement is a law, no way of determining whether a general statement is a law from just the structure of the statement. Compare 7S.8.5 and 7S.8.6.

7S.8.5 All gold spheres are less than one mile in diameter.
7S.8.6 All uranium spheres are less than one mile in diameter.

7S.8.5 and 7S.8.6 are grammatically, syntactically, identical. Yet 7S.8.5 is not a law and 7S.8.6 is a law. We could, if we had certain peculiar political or alchemical abilities, gather together a mile-diameter ball of gold. Such a ball of uranium would exceed its critical mass and explode. The difference is due to the nature of the world, not the nature of the language or logic.

Even if we had a good way to pick out the general laws like 7S.8.6 and to distinguish them from accidental generalizations like 7S.8.5 and the first premise of 7S.8.4, there remains a problem of specifying what an instance is, which particular events or objects a law applies to. Consider 'all ravens are black'. It is natural to think that it applies to all and only ravens. But, given the logical structure of D-N explanations, the domain of objects to which such a law applies might be much broader, as we will see in our next two subsections.

There are many theories of explanation other than the D-N model. Many of them are refinements of the D-N account, emphasizing the causal nature of laws, for example. Insisting that an explanation must focus on a cause may help with the flagpole example: the pole causes the shadow, but the shadow does not cause the pole.

Other important refinements take into account the probabilities involved in many explanations. When we explain an instance of lung cancer, for example, by the fact that someone smoked cigarettes habitually, we appeal to general, but not categorical, laws, about the relationship between smoking and lung cancer. Many laws appear to be probabilistic, as do the fundamental laws of quantum mechanics.

Some philosophers of science appeal to unification as an essential aspect of explanation: a good explanation unifies disparate phenomena. Other philosophers take explanation to be a pragmatic or epistemic notion; rather than explaining particular claims by appeals to general claims, we explain what we don't know in terms of what we do know, explaining the unfamiliar by appeal to the familiar. Such models of explanation invoke the inferential tools of logic less apparently than the core D-N model, though the extent of their reliance on inference varies.

LAWS AND INDUCTION

We saw in the above discussion of explanation that the difference between laws and law-like generalizations is difficult to discern. This problem of determining the laws is important in most accounts of explanation, and in the philosophy of science more generally. One aspect of the problem is determining when an event confirms a general claim. Our studies of formal logic are useful in understanding the nature of confirmation and some problems that arise in refining that notion.

Confirmation, on some accounts, is the converse of explanation, and so raises some similar problems. Where a law and initial conditions might explain an event or phenomena, a particular event can confirm or disconfirm a hypothesis, one that may even turn out to be a law. Like explanation, the problem of confirmation involves specifying the connection between a general claim and a particular statement, often between a theory and an observation.

In science, we generally want to summarize myriad diverse experiences, reducing them (in some sense) to a small set of general principles. For example, we might see an apple fall to the ground. Then another, and another. We can summarize all these individual events, as at 7S.8.7.

7S.8.7 All apples fall to the ground.

7S.8.7 is partly general, in that it applies to all apples, and partly limited, in that it applies only to apples. We might then notice that pears and peaches and oranges also fall to the ground, and develop the more general 7S.8.8.

7S.8.8 All fruit falls to the ground.

7S.8.8 represents an increase in the generality of 7S.8.7, but is still limited to fruit. We might further notice that vegetables and rocks and people also fall to the ground. We could propose 7S.8.9.

7S.8.9 All objects fall to the ground.

7S.8.9 is satisfyingly more general than 7S.8.7 and 7S.8.8. It looks a lot like something we would be happy to call a law of nature.

Unfortunately, as it stands, 7S.8.9 is false, entailing precipitously that smoke and steam and helium balloons are not objects. We thus have to refine 7S.8.9, replacing a rough concept like 'falling to the ground' with concepts like density, force, and gravity. Newton's work on gravity unified accounts of terrestrial and celestial motions, making the laws of motion much more general, since they apply to any two objects, whether on Earth or in the sky, and yields the general law of what is now called Newtonian gravitational theory at 7S.8.10.

7S.8.10 $F = Gm_1m_2/r^2$

7S.8.10 applies to any two objects. (Well, not mathematical objects, or other abstract objects like propositions, but let's not worry about those here.) Still, 7S.8.10 continues to be limited, or anyway an incomplete description of the motion of objects.

To describe fully the motion of a particular object, we have to introduce other forces, ones that act at the same time as gravitational force, like electromagnetic force.

The point here is that we want scientific theories to be general, summary claims, applying to as many specific cases as possible. The process by which we organize a limited number of experiences into a general claim is called induction. The problem of induction is that the leap from the specific to the universal, even in 7S.8.7, involves appeal to claims about causal connections of which we have no experiences. That is, if we want statements like 7S.8.7–7S.8.10 to apply to future and unseen fruit (and other objects), then we need to presume something like a law of the uniformity of nature, some claim that the future and unobserved will be like the present and observed. From our experiences, we can conclude only claims like 7S.8.7′ or 7S.8.10′.

7S.8.7′	All observed apples have fallen to the ground.
7S.8.10′	$F = Gm_1m_2/r^2$, as far as we have observed.

The predictive force of 7S.8.7 and 7S.8.10 is lost in 7S.8.7′ and 7S.8.10′. Without presuming that our experiences of the past will continue into the future, we fail to meet a central demand of science.

One solution to this problem, from the eighteenth century Scottish philosopher David Hume, is to give up any claims about uniformity in nature, and to explain our claims 7S.8.7–7S.8.10 in terms of our expectations. We are built in such a way that we form mental habits, when seeing falling fruit and such, to develop expectations that the future will be like the past. We see apples untethered, and our past experience leads us to believe that the next untethered apples we see will fall, rather than rise or hover. We need not claim insight into the inner workings of nature (of apples or causation) to make this conclusion. We need merely observe that this is the way that our minds work.

One problem with Hume's solution is that it is not always easy to figure out what to expect. We do not always know which claims are confirmed by an experience, and so do not know how to predict novel cases.

Recall the case at 7S.8.4, and imagine that a new student, age twenty, comes into the classroom. This does not increase our confidence in (or serve to confirm) the claim that all persons in the room are under sixty, taken as a predictive law about the classroom. We know that the aged dean, or a visiting parent, or an older professor could just as easily enter the classroom. That a person in a room is under sixty will confirm the hypothesis that all persons in a room are under sixty. But it will not increase our confidence that future people in the room will be under sixty.

Contrast this situation with one in which we discover that a piece of copper conducts electricity. In this case, we are led to believe that the next piece of copper we encounter will conduct electricity. In both cases, we are presented with regularities. Only in some cases are we led (by habit or whatever means) to expect that this regularity will continue to apply, that we have, in summarizing the observed facts, discovered a law that will allow us to predict future events.

The difference between the copper case and the under-sixty case is that one is law-like and the other is not. But to say that one is law-like and the other is not is merely to restate the problem, not to solve it. The question remains how to characterize the difference between law-like and non-law-like generalities. And this is essentially, at least in part, the problem of determining when an experience confirms a general claim and when it does not.

CONFIRMATION

Let's think more about the relationship between theories and the evidence for them. To start, let's return to the classic example of a law-like statement we saw earlier, "All ravens are black." While most of what we think of as scientific laws or theories are more complicated than this simple claim, it is, like many laws, a generalization, plausibly formulated on the basis of induction on observations, and one that projects into the future: not only are the ravens I have seen black, I also expect that the ravens I encounter in the future will be black. Similarly, all objects I encounter in the future will have the gravitational attraction between them that laws of gravitation ascribe, all charged particles will have the electromagnetic forces between them that Coulomb's law describes, all neurological processes will have the characteristics that neuroscience describes, and so on.

When general claims such as these are first developed, they are considered hypotheses, ones that we do not know whether to believe. Questions about whether to believe a hypothesis are questions about confirmation: What evidence increases my belief in the hypothesis? How much more credence should I give to a hypothesis, given some particular event or observation? When, if ever, do we take a hypothesis as fully confirmed?

To take a simple case, and again following Hempel's seminal work, this time on confirmation, when we consider a hypothesis of the form 'all Ps are Qs', we can identify both evidence that confirms the hypothesis and evidence that refutes it. A P that is not a Q—for example, a red raven—decisively refutes the hypothesis. A P that is a Q—a black raven—confirms the hypothesis.

This simple, perhaps obvious, point, which we saw in section 6.2 of *IFLPA*, is known as Nicod's criterion, after the early twentieth-century logician and philosopher Jean Nicod. Nicod's criterion raises at least two kinds of questions. First, how much should our belief in a hypothesis be increased on the evidence of a P that is a Q? Returning again to the case at 7S.8.4, a twenty-year-old student entering my classroom will, according to Nicod's criterion, confirm the claim that all people in the room are under sixty, but not so much as to make us think that the hypothesis is law-like. In contrast, a single piece of copper conducting electricity will lead us to believe that any future piece of copper will conduct electricity.

Whether and to what degree an observation will confirm a hypothesis is thus a complex matter worth detailed consideration. Sometimes a single case suffices; other times, no number of examples will. Formal logic has a significant role in many theories of confirmation. In particular, some views of confirmation take the relation between evidence and a theory to be deductive, just like the relation between explanations and phenomena. On the hypothetico-deductive view of confirmation, a hypothesis is confirmed by some observation or claim if the hypothesis formally entails that claim, given some auxiliary claims, including a logical theory. The complications arise when one tries to spell out the auxiliary claims needed beyond the logic.

The second kind of question raised by Nicod's criterion concerns the relation between claims that all Ps are Qs and our observation of things that are not Ps, the relation, that is, between 'all ravens are black' and white swans or blue bonnets. Naturally, we would like to say that something that is not a P has no relevance to the claim that all Ps are Qs. Blue bonnets are irrelevant to our belief that all ravens are black. But it is also natural to think that anything that confirms a statement will confirm any statement logically equivalent to it. Logically equivalent statements, after all, have the same truth conditions, so should be true or false on the same evidence. This latter claim is sometimes called Hempel's equivalence condition.

Unfortunately, all of these natural beliefs lead us to the paradox of the ravens, so called after the classic example we are considering. 7S.8.11 and 7S.8.12 are logically equivalent, by the rule of contraposition.

7S.8.11 $(\forall x)(Px \supset Qx)$
7S.8.12 $(\forall x)(\sim Qx \supset \sim Px)$

Taking 'Px' as 'x is a raven' and 'Qx' as 'x is black', we are led to think that a black raven confirms 7S.8.11, a non-black raven refutes the claim, and anything that is not a raven has no relevance to the claim.

But on the same interpretation of the predicates, 7S.8.12 says that all non-black things are not ravens. On Nicod's criterion, 7S.8.12 is confirmed by anything that is non-black non-raven, like a white swan or a blue bonnet. Moreover, anything that is black (that is, anything that is not non-black, like a raven) is irrelevant to the claim!

There are many options to avoid the paradox. One could drop Hempel's equivalence condition that observations that confirm a claim also confirm all logically equivalent claims. The consequences for a logic of scientific theory are potentially devastating, though alternative logics are possible.

Alternatively, one could bite the bullet and say that non-Q non-Ps actually do confirm the claim that all Ps are Qs, though perhaps not as much as a P that is Q. One might claim that an observation of a white swan confirms that all ravens are black, but only in proportion to the relevant populations. There are many more non-black things than ravens, so the observation of a non-black thing that is not a raven is not as influential on our beliefs about 7S.8.12 (or about 7S.8.11) as the observation of a raven that is black.

Hempel favored a solution like the latter, noting, with Nelson Goodman, that confirmation takes place within a background of information and beliefs. Understanding the context of a hypothesis, and evidence for or against it, is essential to the development of scientific theory, which is not merely a formal logical matter.

Nevertheless, not only logic, but probability theory has come to be essential to our understanding the nature of confirmation. On what is known as a Bayesian view, after the eighteenth-century statistician Thomas Bayes, evidence confirms a hypothesis, to some degree, if it raises the probability of the hypothesis. Working out what it means to increase the probability of a hypothesis has turned out to be difficult, and has led to fecund work on probability, beliefs, and science.

RESOLVING CONTRADICTIONS

We have looked at two aspects of scientific methodology that may be illuminated by our studies of formal logic: explanation and confirmation. Let's look at one more aspect of scientific method with a connection to logic. This one is less philosophically interesting (by which I mean that it is somewhat less contentious), but it is more closely connected to our work in logic.

We have lots of beliefs. Some of them are explicit, and can be recalled or affirmed at will, as my belief that Abraham Lincoln was the sixteenth president of the United States of America. Others are less easy to state or recall, like my beliefs about mathematics or how to ride a bicycle. The former are difficult to state in part because there are so many of them. The latter seem difficult to state because they are largely implicit.

We manage our beliefs constantly, updating them as we receive new evidence and giving up or restricting beliefs in general claims as contravening evidence appears. The process of managing our belief system is often, perhaps always, guided by principles about what makes a good set of beliefs. What are the claims I should believe, given the evidence available to me? That guiding question should make it clear that epistemology, the study of knowledge and beliefs, is a normative field, not purely a descriptive one: my belief set is formed both by facts about what I believe and principles about what I should believe.

Principles about what I should believe are not really separable from principles of scientific theory construction. Scientific methods are not isolated to lab work in science departments. They are methods of maintaining a healthy set of beliefs. We are all scientists, in our everyday life, and our methods for evaluating our beliefs, if they are to be the best methods, do not differ from the scientific method. In both domains, we are guided by some basic principles about the relation of evidence to theories, by our understanding of how the world works, and by a desire to avoid inconsistency.

Formal logic regiments the principle of avoiding inconsistency in a couple of ways. For one, our classical systems of logic are explosive, as we saw in section 3.5: if a theory contains a contradiction, then every formula is provable. Since we want our theories to be sound, we have to make sure that our logic does not contain contradictions. That's the principle underlying our method of indirect proof, or *reductio ad* *absurdum* argumentation: if a formula leads to a contradiction, it must not be true and thus should not be provable.

One of the most important methodological lessons we take from science concerns how to manage a belief that contradicts ones we already hold. Given our fallibility and that our beliefs are not all explicitly available, we ordinarily hold some contradictory beliefs. But sometimes these contradictions are made explicit and once we find that there is a conflict in our set of beliefs, we are responsible for looking at our evidence and ridding ourselves of the problem.

Ideally, when finding a contradiction among our beliefs, we are faced with a system of hypotheses, each of which is independently justified, but which together are incompatible. We have to choose which hypothesis to cede. A contradiction within a large theory merely tells us that there is a problem in the theory. It need not tell us where the problem lies. So we look at the various evidence.

For example, imagine that we believe that there are going to be no parties this weekend, but then we receive a flyer for a gathering on Friday. Adding the belief we gain from the flyer to our belief set leads to a contradiction. We could resolve that contradiction in various ways, some (but not all) of which are listed at 7S.8.13–7S.8.15.

7S.8.13 We could check the date on the flyer; maybe there is a confusion about the date.

7S.8.14 We could give up our belief about there being no parties this weekend.

7S.8.15 We could redefine the term 'party' such that the gathering is not a party.

More technically, if we have a theory that yields a claim that is inconsistent with new evidence, we only know that assimilating the new claim with the original theory yields some kind of contradiction. A theory is a set of sentences. Let's imagine a theory T, equivalent to the set of sentences at 7S.8.16.

7S.8.16 $S_1 \bullet S_2 \bullet S_3 \bullet \ldots \bullet S_n$

Let's further imagine that T yields some claim O, but that we get new information that contradicts that claim, ~O. By modus tollens, T is false. But T is just the conjunction at 7S.8.16. We are left with the negation of a series of conjuncts at 7S.8.17.

7S.8.17 $\sim(S_1 \bullet S_2 \bullet S_3 \bullet \ldots \bullet S_n)$

7S.8.17 is, by an extended version of De Morgan's law, logically equivalent to the series of disjuncts at 7S.8.18.

7S.8.18 $\sim S_1 \vee \sim S_2 \vee \sim S_3 \vee \ldots \vee \sim S_n$

That's as far as the logic will take us. As a logical matter, we don't know which of the sentences of the theory to reject in order to restore consistency to our believe set. In indirect proofs, we make one explicit assumption and always abandon that one when we reach a conclusion. Reasoning in the real world is messier.

Given that we have various options, we need methods for weighing the evidence, for choosing among those options. Those methods are governed by various abstract principles. We look for underlying beliefs that might be at fault, perhaps bringing some implicit beliefs to our awareness, or reconsidering evidence for them. We can easily restore consistency to a theory that contains a contradiction in different, incompatible ways, so we work carefully to make sure that we pick the best option available.

Scientific theories, like our belief systems, are generally underdetermined by evidence. Given the same experiences, I could have different beliefs. Simple examples include the fact that evidence often provides correlation without indicating causation. For example, a recent study shows that Facebook users get lower grades in college. We do not know whether to conclude that Facebook use causes lower grades or that people who use Facebook are those who are already likely to be less successful. Similarly simple examples are ubiquitous.

More profoundly, we have choices among theories for holistic reasons. When presented with a theory and an observation, there are various options of how to integrate the observation into the theory. Even when an observation does not conflict with other, previously accepted hypotheses, there are always lots of theories that can accord with our claims. Some of those theories have extraneous elements. We discount theories that refer to ghosts, for example, and seek an explanation of the noise in the attic that appeals only to natural phenomena, like wind and expansion or contraction of materials due to humidity. We invoke principles of parsimony, or Ockham's razor: do not multiply entities beyond necessity.

There are various virtues that guide our analyses of hypotheses, that help us to determine how best to restore consistency to our theory, or belief set, and that govern scientific reasoning generally. Among them are properties including modesty, simplicity, and generality. We accept only the weakest, or most modest, principles as the most plausible. We generally view the simplest explanations of phenomena as the most likely, though simplicity is not a categorical criterion: there are different ways for theories or claims to be simple, and they are sometimes in tension with one another. The claim that objects fall to Earth is simple, but it conflicts with gravitational theory, which is better because it is more general.

Science is the epitome of a rational enterprise, especially when we consider science in its broadest form. Our methodological principles are not merely for scientists in lab coats. They are guiding principles for managing our beliefs, including for philosophical theorizing. All reasoning is governed by the same kinds of principles and the same scientific methods, variously instantiated and applied.

TELL ME MORE

- What is Nicod's criterion, and how is it important for conditionals? See 6.2: Conditionals.
- What is the problem of induction? How does it undermine attempts to identify laws? See 7.1: Deduction and Induction.

For Further Research and Writing

1. What is scientific method? How does it relate to our ordinary reasoning? What is the role of formal logic in managing our beliefs and our scientific hypotheses? The readings from Quine and Ullian and from Papineau will be useful.
2. What is the role of logical entailment in an explanation? Describe the deductive-nomological model of explanation and at least one objection to it. Hempel's *Philosophy of Natural Science* will be especially helpful here.
3. What are the alternatives to the deductive-nomological model of explanation? Compare and contrast at least one other approach. You might find the Toulmin, Kitcher, Friedman, or Van Fraassen readings useful, in addition to Hempel's work.
4. What is the relation between causation and explanation? Do the logical entailments of the D-N model capture the causal relations? See especially Salmon's work.
5. What is a probabilistic explanation? Can the deductive-nomological model of explanation be adapted to provide one? See Hempel's *Philosophy of Natural Science* and Railton's article.
6. What does it mean to confirm a claim or a theory? Is the best notion of confirmation a logical relation? See Huber for a good overview and Hempel's paper on confirmation for the formal-logical approach.
7. How do considerations of the problems of induction (see section 7.1) affect questions about explanation or confirmation? In addition to Hume and Goodman, you might find the Huber and Papineau pieces useful.
8. What is the paradox of the ravens? What is the role of formal logic in formulating the problem? How might it best be addressed? See Papineau's introduction, Hempel's papers on confirmation, and Goodman's few interesting comments in "The New Riddle of Induction" in *Fact, Fiction, and Forecast*.

Suggested Readings

Fetzer, James. "Carl Hempel." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2015. http://plato.stanford.edu/entries/hempel/. An excellent overview of the work of one of the most important philosophers of science.

Friedman, Michael. "Explanation and Scientific Understanding." *Journal of Philosophy* 71, no. 1 (1974): 5–19. Friedman evaluates both the D-N model and other approaches that emphasize understanding, seeking an account in terms of unification.

Goodman, Nelson. *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press, 1955. Goodman's articles here on counterfactual conditionals and the new riddle of induction are especially important, and include discussion of the paradox of the ravens as well.

Hempel, Carl Gustav. *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press, 1965. This collection contains Hempel's papers on explanation and confirmation, and other important essays.

Hempel, Carl Gustav. *Philosophy of Natural Science*. Prentice Hall 1966. An excellent and accessible introduction to the philosophy of science from one of the most important philosophers of science.

Hempel, Carl Gustav. "Studies in the Logic of Confirmation." *Mind* 54 (1945): 1–26, 97–121. This essay, in two parts, is a detailed examination of the problems facing a logic of confirmation.

Hempel, Carl Gustav, and Peter Oppenheim. "Studies in the Logic of Explanation." *Philosophy of Science* 15 (1948): 135–175. The classic Hempel paper on explanation.

Huber, Franz. "Confirmation and Induction." In *The Internet Encyclopedia of Philosophy*, edited by James Fieser and Bradley Dowden. http://www.iep.utm.edu/conf-ind/. A broad survey of the problems of confirmation, including some important topics not covered here and an emphasis on probabilistic reasoning.

Hume, David. *An Enquiry Concerning Human Understanding*. Hume's discussion of the problem of induction sets the stage for all future work in the philosophy of science. See especially sections 4–7.

Kitcher, Philip. "Explanatory Unification." *Philosophy of Science* 48, no. 4 (1981): 507–531. Kitcher pursues Friedman's suggestion that explanation is unification in greater detail.

Losee, John. *A Historical Introduction to the Philosophy of Science*, 4th ed. Oxford: Oxford University Press, 2001. An excellent introductory text to the philosophy of science, with an emphasis on historical debates.

Papineau, David. "Methodology: The Elements of the Philosophy of Science." In *Philosophy 1: A Guide Through the Subject*, ed. A. C. Grayling, 123–180. Oxford: Oxford University Press, 1998. A broad and useful introductory essay on a range of topics in the philosophy of science.

Quine, W. V., and J. S. Ullian. *The Web of Belief*, 2nd ed. New York: McGraw Hill, 1978. The discussion of a murder mystery in the early pages is a friendly discussion of the principles of theory construction and revision.

Railton, Peter. "A Deductive-Nomological Model of Probabilistic Explanation." *Philosophy of Science* 45 (1978): 206–226. The classic paper on probabilistic explanation within a broadly D-N framework.

Salmon, Wesley. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press, 1984. Salmon emphasizes the causal aspects of explanation.

Salmon, Wesley. "Scientific Explanation: Three Basic Conceptions." *Proceedings of the Biennial Meeting of the Philosophy of Science Association* 2 (1984): 293–305. Epistemic, modal, and ontic conceptions of explanation.

Toulmin, Stephen. *Foresight and Understanding: An Enquiry into the Aims of Science*. New York: Harper and Row, 1961. A wonderfully readable overview of some general principles of scientific reasoning, emphasizing human understanding in contrast to Hempel's focus on logical entailment.

Van Fraassen, Bas. "The Pragmatics of Explanation." *American Philosophical* Quarterly 14 (1977): 143–150. Van Fraassen argues that explanation is not a logical or semantic phenomenon, but a human and pragmatic one.