1

What Is Induction and Why Study It?

Evan Heit

Why study induction, and indeed, why should there be a whole book devoted to the study of induction? The first reason is that inductive reasoning corresponds to probabilistic, uncertain, approximate reasoning, and as such, it corresponds to everyday reasoning. On a daily basis we draw inferences such as how a person will probably act, what the weather will probably be like, and how a meal will probably taste, and these are typical inductive inferences. So if researchers want to study a form of reasoning that is actually a pervasive cognitive activity, then induction is of appropriate interest.

The second reason to study induction is that it is a multifaceted cognitive activity. It can be studied by asking young children simple questions involving cartoon pictures, or it can be studied by giving adults a variety of complex verbal arguments and asking them to make probability judgments. Although induction itself is uncertain by nature, there is still a rich, and interesting, set of regularities associated with induction, and researchers are still discovering new phenomena.

Third, induction is related to, and it could be argued is central to, a number of other cognitive activities, including categorization, similarity judgment, probability judgment, and decision making. For example, much of the study of induction has been concerned with category-based induction, such as inferring that your next door neighbor sleeps on the basis that your neighbor is a human animal, even if you have never seen your neighbor sleeping. And as will be seen, similarity and induction are very closely related, many accounts of induction using similarity as their main currency (Heit & Hayes, 2005).

Finally, the study of induction has the potential to be theoretically revealing. Because so much of people's reasoning is actually inductive reasoning, and because there is such a rich data set associated with induction, and because induction is related to other central cognitive activities, it is possible to find out a lot about not only reasoning but cognition more generally by studying induction.

Induction is traditionally contrasted with deduction, which is concerned with drawing logically valid conclusions that must follow from a set of premises. The following section will consider possible definitions of induction by describing possible relations between induction and deduction. But first it is useful to briefly mention that the reasons for studying induction to some extent are linked to the differences between induction and deduction. That is, it could be argued that induction, in comparison to deduction, characterizes more of everyday reasoning, has the potential to be studied with a broader range of tasks and materials, and is closely related to other cognitive activities that help people manage uncertainty.

HOW IS INDUCTION RELATED TO DEDUCTION?

Although it might be natural to ask "how are induction and deduction different?" that would presuppose the conclusion that they are actually different. Although induction and deduction are traditionally considered alternatives to each other, as will be seen under some conceptions the similarities are much greater than the differences. Before assessing to what extent induction and deduction are similar or different, it is first important to consider just what kind of entities induction and deduction are. Although not always made explicit by researchers, there are two views on this issue, namely, the "problem view" and the "process view." According to the problem view, induction and deduction refer to particular types of reasoning problems. So from looking at a particular problem, say a question on a piece of paper in a psychological experiment on reasoning, it should be possible to say whether this is an induction problem or a deduction problem (or possibly it could be deemed debatable whether it is one or the other). In contrast, according to the process view, the locus of the question is not on the paper but in the head. That is, induction and deduction refer to psychological processes. For a given problem, it may be possible to answer it using induction processes or deduction processes. Likewise, we can investigate what is the relation between the two kinds of processing.

The problem view and the process view have to a large extent been confounded in the literature. That is, researchers who have studied problems that are traditionally thought of as induction have typically been interested in different psychological theories than researchers who have studied traditional deduction problems. However, for the sake of clarity it is better to treat the two views separately, namely, how problems of induction may differ from problems of deduction, and how inductive processes may differ (or not differ) from deductive processes. These two views will now be addressed in turn.

The Problem View

GENERAL AND SPECIFIC. It is sometimes said that induction goes from the specific to the general, and deduction goes from the general to the specific. For example, after observing that many individual dogs bark, one might induce a more general belief that all dogs bark. Alternately, having the general belief that all dogs bark, one might deduce that some particular dog will bark. However, there are difficulties with this version of the problem view. Consider the following arguments. (The statement above the line is a premise that is assumed to be true, and the task is to consider the strength of the conclusion, below the line.)

Dogs have hearts (1)

All mammals have hearts

All mammals have hearts (2)

Dogs have hearts

Dogs have hearts (3)

Wolves have hearts

Dogs have hearts (4)

At least some mammals have hearts

Argument (1) is a good example of an inductive argument going from specific to general, and likewise argument (2) is a good example of a deductive argument going from general to specific. Yet arguments (3) and (4) do not fit neatly into this scheme. Argument (3) is somewhat plausible but surely not deductively valid, so it is better thought of as an inductive argument. Yet it goes from specific to specific rather than specific to general. Finally, argument (4) seems to be deductively valid, yet it starts with a specific statement. Still, it is possible to disagree about these last two arguments. For argument (3), it could be said that there is an intervening general conclusion, such as “All mammals have hearts.” For argument (4), it could be said that there is a hidden general premise, such as “All dogs are mammals.” The key point is that one can’t just look at the written form of an argument, in terms of whether it goes from specific to general or vice versa, and easily state whether it is inductive or deductive in nature.

DEFINING VALIDITY. Hence, it would seem profitable to take a more subtle approach to the problem view. Perhaps the most defensible version of the problem view is to define deductively valid arguments, and relate other kinds of arguments to those that are deductively valid. One standard definition of deductively valid arguments is that these are arguments following the rules of a well-specified logic. Assuming that one can specify the rules of one's preferred logic, say in terms of truth tables for various symbolic combinations, then it should be possible (if not easy) to determine whether any given argument is deductively valid or not. It might be seen as a small disadvantage of this approach that deductive validity is not defined in absolute terms but only relative to a logic. Different people might endorse different logics and hence disagree about which arguments are deductively valid. On the other hand, defining deductive validity relative to a logic could be seen as an advantage in terms of giving flexibility and in terms of appreciating that there is not a single true logic that is universally agreed.

A more serious problem with this version of the problem view is that it does not say much about inductive problems. Once the deductively valid arguments are defined, what remains are the deductively invalid ones. Presumably some of these are stronger than others, in terms of induction. For example, compare argument (1) above to argument (5) below.

Dogs have hearts (5)

All living things have hearts

It should be clear that neither (1) nor (5) is deductively valid, yet somehow (1) seems more plausible in terms of being a good inductive argument. Whatever rules of logic are used to define deductive arguments may not be too useful in determining that (1) is stronger than (5).

LEVELS OF CERTAINTY. Another approach to the problem view is to describe arguments in terms of the certainty of their conclusions (Skyrms, 2000). Consider argument (6).

Dogs have hearts (6)

Dogs have hearts

In this case, it seems absolutely certain that if the premise is taken to be true, then the conclusion must necessarily follow. This must be a perfectly valid argument. On the other hand, an argument such as (2) above might seem to have a very certain conclusion, perhaps 99.5% certain. This level of certainty could still be well over the threshold that is required for saying that an argument is deductively valid. Let's say, hypothetically, that arguments with conclusions below the 99% level of certainty will be called deductively invalid. Even among these arguments, this version of the problem view allows a great deal of differentiation. For example, argument (1) might be associated with 80% certainty and argument (5) might be associated with 10% certainty. Hence (1) would be considered a much stronger inductive argument in comparison to (5).

Perhaps the greatest appeal of this version of the problem view is that it allows for deduction and induction to be placed on a common scale of argument strength. In principle, any argument would have a place on this scale, and whether it is deductively valid, inductively strong, or inductively weak would be determined by the value on the scale. The most obvious problem, though, is that there is still a need for assessing the place of each argument on the scale. One nice idea might be an inductive logic, that is, some set of rules or operations that for a set of premises can assign a certainty value for a conclusion.

A subtler problem is that "certainty" itself would need to be defined better. For example, in argument (1), either the conclusion that all mammals have hearts is true or it is not, so the conversion from probability to certainty may not be obvious. For example, it would seem a little funny to assign a certainty level from 0% to 100% to a statement that is either true or false. (Perhaps certainty could be related to the proportion of mammals with hearts, rather than to whether it is true that all mammals have hearts.) Another issue to clarify is the distinction between argument strength and certainty of the conclusion. Argument (1) may seem strong simply because people believe the conclusion that all mammals have hearts. Now compare that argument to argument (7), below.

Lemons have seeds (7)

All mammals have hearts

Here is a situation where the two arguments have the same conclusion, which is equally certain in each case, but (1) seems much stronger than (7). It could be valuable to consider other ways of representing argument strength here, such as the conditional probability of the conclusion given the premise, or the difference between the unconditional probability of the conclusion and the conditional probability of the conclusion, given the premise.

MATTERS OF CONVENTION. A final perspective on the problem view is simply to be descriptive, that is, to enumerate what kinds of arguments are studied by researchers under the topics of induction and deduction. Induction is potentially a very broad topic and a variety of cognitive activities have been referred to as induction, including categorization, reasoning by analogy, and probability judgment. However, many of the chapters in this book focus on a more particular kind of induction, namely, category-based induction, involving arguments about categories and their properties. (Most of the examples in this chapter represent typical examples of category-based induction.) Research on adults' reasoning usually involves presenting arguments like these in written form; however, for children it is possible to present similar information with pictures. Studies of induction typically ask people to make judgments of argument strength, such as to judge which of two arguments is stronger, or with a single argument to make a continuous judgment of strength or probability.

In comparison to induction, research in deduction has used a narrower range of problems. One typical area of research within deduction is conditional reasoning – arguments involving ifs and thens, examining reasoning involving classic rules such as modus ponens and modus tollens. Another area of research within deduction is syllogistic reasoning – reasoning with arguments with statements like "All artists are beekeepers." Indeed, for syllogisms involving two premises, there are only sixty-four classical forms of syllogism. Research on deduction tends to ask people to assess logical validity of conclusions (a yes or no question) rather than make continuous judgments. Overall, the conventional approach is like other approaches to the problem view in that there is a relatively narrow range of arguments corresponding to deduction and a wider, somewhat ill-defined, range corresponding to induction. Yet even within the area of deduction research, there are lively debates about what exactly is a problem of deduction. For example, Wason's selection task, involving selecting cards to test a rule such as "If a card has a vowel on one side then it has an even number on the other side," has been variously argued to be a problem of deduction or induction (e.g., Feeney & Handley, 2000; Oaksford & Chater, 1994; Poletiek, 2001).

EVALUATION OF THE PROBLEM VIEW. Perhaps the most appealing aspect of the problem view is that it offers the possibility of defining induction and deduction in an objective way, in terms of the problem being solved or the question being asked. (The problem view is more impressive in terms of defining deduction in comparison to defining induction, though.) From the point of view of psychologists, this strength would also be the greatest weakness, namely, that the problem view does not itself refer to psychological processes. Just because one problem is defined as an induction problem and another is defined as a deduction problem does not guarantee that people will engage in inductive reasoning processes for one task and deductive reasoning processes for the other task. The same processes could be used for both, or the delimitation between types of psychological processes might not correspond at all to the agreed definition of problems, or any problem might engage a mixture of processes. In the terminology of memory research, there are no process-pure tasks. Of course, for computer scientists or logicians, reference to psychological processes may not be a priority. Still, it does seem desirable to consider the alternative of treating induction and deduction as possible kinds of psychological process. Hence, this chapter will next turn to the process view.

The Process View

According to the process view, comparing induction and deduction is a matter of specifying the underlying psychological processes. According to one-process accounts, the same kind of processing underlies both induction and deduction. Another way to describe this idea is that there is essentially one kind of reasoning, which may be applied to a variety of problems that could be considered either inductive or deductive in nature (Harman, 1999). In contrast, according to two-process accounts, there are two distinct kinds of reasoning. It is possible that these two kinds of reasoning directly correspond to induction and deduction. Alternately, the two kinds of reasoning might correspond to some other distinction, such as intuitive reasoning versus deliberative reasoning, that could be related to the distinction between induction and deduction. It should be acknowledged at the start that one-process and two-process accounts are somewhat poorly named. That is, at some level, reasoning surely involves many different psychological processes. The question, though, is whether the same processing account is applied to both induction and deduction, or whether two different processing accounts are applied. Some examples of one-process and two-process accounts will now be described, followed by the presentation of some experimental evidence aimed at assessing these accounts.

ONE-PROCESS ACCOUNTS. One of the most widely known theories of reasoning is mental model theory, which proposes that people solve reasoning problems extensionally by constructing models of possible states of the world and performing operations and manipulation on them (Johnson-Laird, 1983). Mental model theory is usually thought of as an account of deduction, and indeed it has been extensively applied to conditional-reasoning and syllogistic-reasoning problems. However, it has also been argued that mental model theory can be applied to problems of induction, namely, probabilistic reasoning tasks (Johnson-Laird, 1994; Johnson-Laird, Legrenzi, Girotto, Legrenzi, & Caverni, 1999). Hence, mental model theory is a one-process account, in the sense that it is aimed at giving a singular account for problems both of induction and deduction.

A newer alternative to mental model theory is the probabilistic account, which aims to account for a variety of reasoning phenomena, particularly traditional deduction problems in terms of probabilistic formulas, such as from Bayesian statistics (Chater & Oaksford, 2000; Oaksford & Chater, 1994). Essentially, the probabilistic account is saying that people solve deduction problems by means of induction processes. This account does not propose different kinds of processing for performing deduction, and hence the probabilistic account is also a one-process account.

The previous accounts are aimed mainly at problems of deduction. In contrast, other reasoning accounts have focused on problems of induction, such as category-based induction (Heit, 1998; Osherson, Smith, Wilkie, Lopez, & Shafir, 1990; Sloman, 1993). These accounts are aimed at predicting the judged strength of various inductive arguments, for example, that (1) above seems stronger or more plausible than (5). Typically, these accounts of induction are based on some measure of similarity or overlap between premise and conclusion categories, in terms of existing knowledge. In this example, there is more known overlap between dogs and mammals than between dogs and living things; hence the argument relating dogs and mammals seems stronger than the argument relating dogs and living things. Now refer back to argument (6). Here, there is perfect overlap between the premise category and the conclusion category – in this case the categories are both *dog*. Hence, there is perfect overlap between premise and conclusion categories, and these accounts of induction should predict that (6) is perfectly strong. In other words, accounts of induction can treat some deductively valid arguments as a special case rather than as being wholly different than inductively weak or strong arguments. The same processing mechanisms – for example, for assessing overlap – would be applied to problems of induction and deduction. In this way, these accounts of induction are one-process accounts. However, it should be made clear that these accounts of induction do not give complete accounts of deductive phenomena. For example, many deductively valid arguments in conditional and syllogistic reasoning could not be assessed simply in terms of feature overlap between premise and conclusion categories.

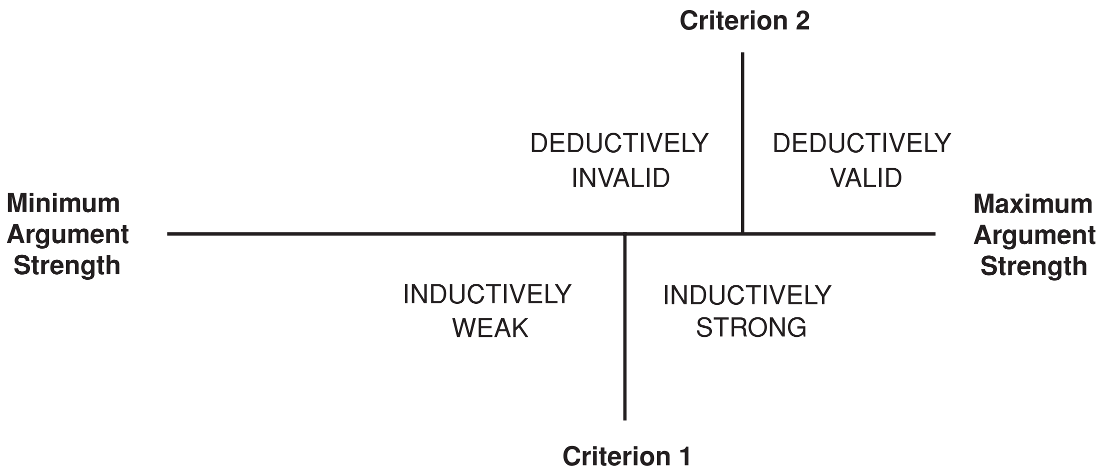

FIGURE 1.1. Criterion-shift account of deduction and induction.

In addition to these one-process accounts that specialize mainly in either deduction or induction problems, there is an alternative that does not give a detailed account of either deduction or induction but does offer a balanced view of how deduction and induction are related to each other. The criterion-shift account (described by Rips, 2001) is closely related to the levels-of-certainty version of the problem view and is illustrated in Figure 1.1. Under this account, assessing the strength of an argument involves finding its place on a one-dimensional scale ranging from minimum argument strength (the most unconvincing argument possible) to maximum strength (an utterly and completely compelling argument). To assess whether an argument should be considered inductively strong, its strength is compared to a criterion, such as criterion 1 in the figure. To assess whether an argument is deductively valid, the criterion is shifted to the right, to criterion 2. By this criterion, an argument would have to be judged extremely strong before it could be called deductively valid. The same reasoning mechanisms would be used for different argument types. The only difference between performing induction or deduction would be represented as a shift of the criterion.

TWO-PROCESS ACCOUNTS. In contrast to one-process accounts, other researchers have emphasized a distinction between two kinds of reasoning (e.g., Evans & Over, 1996; Sloman, 1996; Stanovich, 1999). In these two-process accounts there is one system that is relatively fast but heavily influenced by context and associations, and another system that is more deliberative and analytic or rule based. These two systems do not necessarily correspond directly to induction and deduction. That is, the traditional distinction between these two forms of reasoning may not be the best way to divide things in psychological terms. Still, it is plausible that induction would depend more on the first system, whereas deduction would depend more on the second system. These two-process accounts have been used to explain a variety of findings in reasoning, concerning individual differences, developmental patterns, and relations between reasoning and processing time. For example, in Stanovich's work there is a rich account of how reasoning in the more deliberative system is correlated with IQ, accounting for patterns of individual differences in a variety of problems that would rely more on one system or the other.

EVALUATING ONE-PROCESS AND TWO-PROCESS ACCOUNTS. How would it be possible to decide in favor of either one-process or two-process accounts? There is some neuropsychological evidence, based on brain imaging, for two anatomically separate systems of reasoning (Goel, Gold, Kapur, & Houle, 1997; Osherson et al., 1998). In the studies, subjects were given a set of arguments to evaluate. Half the subjects were asked to judge deductive validity and the other half were asked to judge inductive plausibility. Within each study, there were distinct brain areas implicated for deduction versus induction. What is particularly relevant for present purposes is that the problems were the same for the two conditions, but subjects were asked to perform either deduction or induction. Hence, this is a case of unconfounding the process view from the problem view – presumably all that varied between conditions was processes, unlike the situation in most previous studies of deduction, induction, or both, which used very different problems for one task or the other.

One does not require expensive brain imaging equipment to compare deduction versus induction instructions for a common set of problems. Rips (2001) used the same logic in a much cheaper pen-and-paper study, in which subjects were instructed to judge either deductive correctness or inductive strength for two types of arguments. One type of argument was deductively correct but causally inconsistent, such as "Jill rolls in the mud and Jill gets clean, therefore Jill rolls in the mud," and the other type was deductively incorrect but causally consistent, such as "Jill rolls in the mud, therefore Jill rolls in the mud and Jill gets dirty." In terms of the criterion-shift version of the one-process account, if one type of argument is stronger than another for deduction, then the same type of argument should be stronger for induction. In Figure 1.1, let the dots on the scale represent different types of argument. If one type is stronger, that is, further to the right end of the scale, then it should be stronger regardless of whether the induction or deduction criterion is used. Yet the results were that subjects in the deduction condition gave more positive judgments to the correct but inconsistent arguments, whereas subjects in the induction condition gave more positive judgments to the incorrect but consistent arguments. Rips concluded that these results contradicted the criterion-shift account, which predicted a monotonic ordering of arguments in the two conditions. (See Heit & Rotello, 2005, for further examinations of this kind, leading to the same conclusion.)

In sum, there is some evidence already that giving people deduction versus induction instructions leads to qualitatively different results. It would seem difficult to explain these results by assuming that deduction and induction processes are essentially the same, except that deduction involves a stricter criterion for giving a positive response. Yet at the same time, it seems too early to abandon the one-process accounts, which do provide detailed and accurate predictions about a range of phenomena, usually either concerning deductive or inductive problems. In contrast, the two-process view does not seem to be as well developed in terms of providing detailed process accounts and predictions. More experimentation, directly aimed at comparing the one- and two-process views and at further developing the two-process view, is clearly needed.

At a more general level, the process view itself seems to be a rich and worthwhile approach to studying induction. Certainly for psychologists, the problem view does not seem viable. It is a mistake to assume that people are performing deduction processes on designated deduction problems and induction processes on designated induction problems. Likewise, even for psychologists who are developing process level accounts of reasoning, it is important to keep in mind the wide range of possible reasoning problems. Assuming there is at least considerable overlap between deduction and induction processes, an ideal theory of reasoning would not be limited to either traditional problems of deduction or induction but would encompass both types of problem.

APPROACHES TO STUDYING INDUCTION: THE CASE OF THE DIVERSITY PRINCIPLE

Now that the position of induction with respect to deduction has been explored, the next step will be to introduce the study of induction in more depth. Of course, this whole book is devoted to presenting research on induction, so the following material is intended to serve as a microcosm of the book rather than a complete review of induction research. The focus of the remainder of this chapter will be the diversity principle, namely, the idea that more diverse evidence seems to be stronger evidence than less diverse evidence. This principle seems to be ubiquitous in various forms of reasoning. For example, in social reasoning, seeing a person act in an extroverted way, in a variety of contexts, suggests that the person is truly extroverted by nature. In comparison, seeing a person repeatedly act in an extroverted way, but only in one context, does not give as good support for inferences that the person is truly extroverted. In legal reasoning, having two independent witnesses to a crime seems to be stronger evidence than two witnesses who may share a similar perspective, such as two members of the same family. Note that none of this represents anything like a deductively valid inference; however, there is still a sense of stronger evidence coming from more diverse observations.

The diversity principle has been an object of interest in terms of various approaches to induction, including the historical approach, the philosophical approach, the experimental approach, the developmental approach, and the model-based approach. These approaches will be reviewed in relation to how they have addressed the diversity principle.

The Historical Approach

The diversity principle has been historically important in science. Scientists have tended to favor testing a theory with a diverse set of experiments rather than repeatedly conducting the same experiment or close replications. Imagine two scientists who are testing their respective theories. One scientist essentially conducts the same experiment ten times, whereas the other conducts a variety of experiments assessing different aspects of the theory. Which theory will seem to have stronger support?

An early example of the diversity principle in action was in Bacon's *Novum Organum* (1620/1898), which cautioned scientists of the day against inferences drawn from narrow samples. Bacon illustrated this point with the concept of heat, listing twenty-eight different kinds of heat and hot things that would need to be observed in a study of heat. For example, Bacon listed the rays of the sun, fiery meteors, natural warm baths, hot vapor in furnaces, sparks, stacks of hay, and the insides of animals.

In a somewhat more modern example, Salmon (1984) described how early in the twentieth century, chemists and physicists had developed a wide variety of experimental methods for deriving Avogadro's number (6.02×10^{23}), the number of particles in a mole of any substance. These methods included Brownian movement, alpha-particle decay, X-ray diffraction, black-body radiation, and electrochemistry. Together, these techniques led to strong support for the existence of atoms and molecules because the existence of these particles was the basis of a unified explanation for a set of highly diverse results. Salmon argued that any one of these techniques taken alone, no matter how carefully applied, would not have been sufficient to convince scientists of that period to accept the atomic theory over its principal rival, known as energeticism, which conceived of matter as being continuous rather than being composed of particles.

More generally, it should be possible to study inductive reasoning by studying historical examples of reasoning, whether by scientists or others. One advantage of studying scientific reasoning is that the evidence and theories are usually explicit, in comparison to just studying everyday examples of reasoning. (See the previous book on induction, by Holland, Holyoak, Nisbett, and Thagard, 1986, as well as a more recent collection by Gorman, Tweney, Gooding, and Kincannon, 2005, for further examples.) There is an interesting parallel between the historical approach to the study of induction and the historical approach to the study of creativity (e.g., Weisberg, 1986). In each case, it seems that much can be learned about cognition by looking at paragon cases of thinking and reasoning, even outside the bounds of tightly controlled psychological studies.

The Philosophical/Normative Approach

This historical approach is complemented by attempts of philosophers and statisticians to argue that certain patterns of induction are normative or correct. The philosophical approach follows in the footsteps of Hume's (1777) problem of induction, namely, the problem of whether an inductive inference can ever be justified or considered valid. Without solving Hume's problem in absolute terms, it still might be possible to argue that performing induction in one way is better than another. As such, there have been various attempts to argue for or even prove the greater strength of diverse evidence in comparison to less diverse evidence.

For example, Nagel (1939) argued that a scientific theory should be derived from diverse observations to obtain more reliable estimates. He gave the example of inspecting the quality of coffee beans delivered on a ship. It would be better to inspect small samples of beans from various parts of the ship than to inspect a large number of beans from just one location. Carnap (1950) linked the collection of diverse evidence to the principle that scientific theories should make novel predictions rather than merely redescribe old data. Similarly, Hempel (1966) related the collection of diverse evidence to a falsifying research strategy, namely, it is better to test theories with a wide variety of challenging experiments.

These intuitive arguments did not lead directly to more formal proofs of the diversity principle. For example, Carnap (1950) promised a proof of the diversity principle in a future edition of his book that never did appear. More recently, however, there have been several attempts to formalize the advantage for following the diversity principle. As reviewed by Wayne (1995), there have been two approaches. The first approach compares correlated sources of evidence to independent sources of evidence, in statistical terms. For formal treatments of this correlation approach, linking similarity to probability theory, see Earman (1992) and Howson and Urbach (1993). The second approach is the eliminative approach. The idea behind the eliminative approach is that diverse data sets will be particularly useful for eliminating plausible but incorrect hypotheses, allowing stronger inferences to be drawn based on the remaining, contending hypotheses. In contrast, non-diverse data sets will likely be consistent with too many hypotheses to allow any strong inferences. For a formal treatment of this approach, including a geometric proof, see Horwich (1982), and see Heit (1998) and Tenenbaum and Griffiths (2001) for some psychological applications.

Moreover, there have been arguments that following the diversity principle is not normative. For example, using Earman's (1992) derivations of the diversity principle, Wayne (1995) showed that there can be exceptions, namely, that non-diverse observations can lead to strong inferences if this evidence is nonetheless very surprising. Wayne pointed to the case of the near-simultaneous discovery in 1974 of a previously unknown subatomic particle in two laboratories being a case of non-diverse evidence that still had strong implications for revision of theories in physics. Lo, Sides, Rozelle, and Osherson (2002) raised a related criticism of the normative status of the diversity principle. They too argued that what is crucial is not diversity of observations but rather surprisingness of observations. Lo et al. also suggested a set of exceptions, such as the following:

Squirrels can scratch through Bortex fabric in less than 10 seconds. (8)
Bears can scratch through Bortex fabric in less than 10 seconds.

All forest mammals can scratch through Bortex fabric in less than 10 seconds.

Squirrels can scratch through Bortex fabric in less than 10 seconds. (9)
Mice can scratch through Bortex fabric in less than 10 seconds.

All forest mammals can scratch through Bortex fabric in less than 10 seconds.

It seems intuitive that squirrels and bears are a more diverse pair than squirrels and mice. Yet Lo et al. argued that (9) is stronger than (8), because the evidence about squirrels and mice is more surprising than the evidence about squirrels and bears. That is, the knowledge that small animals are less capable of feats of strength than are large animals makes the evidence about squirrels and mice more surprising than evidence about squirrels and bears.

Heit, Hahn, and Feeney (2005) argued that these exceptions to the diversity principle, suggested by Wayne (1995) and Lo et al. (2002), are indeed exceptions, but they do not undermine the normative status of the diversity principle itself. In the example of the discovery of a new subatomic particle in 1974, physicists were influenced not only by diversity but also by many other sources of knowledge in particle physics. In the example of scratching through Bortex fabric, people would be influenced not only by diversity but also by other knowledge about animals and their strength. In other words, these exceptions as stated do not contain all the premises upon which the arguments are based. Reasoning about these arguments is also influenced by other hidden premises or background knowledge, so that diversity is not being assessed in isolation. Therefore, these counterexamples do not invalidate the diversity principle, because they are not pure tests of diversity. Rather, they show that people will use other knowledge when possible. Indeed, philosophers of science have not claimed that the diversity principle is the sole principle for assessing evidence. For example, Popper (1963, p. 232) listed diversity of supporting evidence as one of six criteria for assessing a scientific theory.

In more general terms, it should be possible to consider a variety of patterns in inductive reasoning in the light of the normative question, namely, what is good reasoning. For example, one of the most pervasive findings in psychological research on induction is the similarity effect, namely, that arguments such as (3) above concerning dogs and wolves are considered stronger than arguments such as (10).

Dogs have hearts. (10)

Bees have hearts.

Argument (10) seems much weaker, and this seems to be a consequence of the low similarity between dogs and bees. However, providing a proof of why it is normative to reason on the basis of similarity has been an elusive task for philosophers and psychologists (Goodman, 1972, but see Shepard, 1987).

The Experimental Approach

In addition to the normative perspective on the diversity principle, there has been a sustained effort by psychologists to document how well the diversity principle serves as a descriptive account of how people carry out informal, inductive reasoning. Osherson et al. (1990) documented diversity effects in adults by using written arguments like the following:

Hippos require Vitamin K for the liver to function. (11)
Rhinos require Vitamin K for the liver to function.

All mammals require Vitamin K for the liver to function.

Hippos require Vitamin K for the liver to function. (12)
Hamsters require Vitamin K for the liver to function.

All mammals require Vitamin K for the liver to function.

The subjects judged arguments like (12) to be stronger than arguments like (11), in response to the greater diversity of hippos and hamsters compared to hippos and rhinos. Indeed, there is a great deal of evidence that adults, mainly Western university students, follow the diversity principle when evaluating written arguments (see Heit, 2000, for a review).

However, when looking to other subject populations, and to evidence collected at a greater distance from the psychology lab, there seem to be exceptions to the diversity principle as a descriptive account. In their study of Itzaj-Mayan adults from the rainforests of Guatemala, Lopez, Atran, Coley, Medin, and Smith (1997) did not find evidence for diversity-based reasoning, using arguments with various living things and questions about disease transmission. Indeed, sometimes the Itzaj reliably chose arguments with non-diverse premise categories over arguments with diverse categories. It appears that they were using other knowledge about disease transmission that conflicted with diversity-based reasoning. For example, given a non-diverse argument that two similar kinds of tall palm trees get some disease, one person claimed it would be easy for shorter trees, located below, to get the disease as well. Giving further support to this idea that other strategies and knowledge can overrule diversity, Proffitt, Coley, and Medin (2000) reported that American adults who are tree experts (such as landscapers) did not show strong diversity effects when reasoning about trees and their diseases. The tree experts seemed to be relying on the knowledge that tree diseases tend to spread readily within tree families such as elms and maples.

Medin, Coley, Storms, and Hayes (2003) reported further exceptions to the diversity principle. One exception, referred to as non-diversity by property reinforcement, potentially makes a direct challenge to the diversity principle that is not as easily explained in terms of the use of other knowledge. The idea behind non-diversity by property reinforcement is that two diverse categories may nonetheless have some characteristic in common and tend to generalize only to other categories with this same characteristic. In the non-diversity by property reinforcement effect, "if an otherwise diverse set of premises shares a salient property not shared by the conclusion category, the reinforcement of the property might weaken that argument relative to a related argument with less diverse premises" (p. 523). This phenomenon is illustrated by the following example:

Polar bears have property X. (13)
Antelopes have property X.

All animals have property X.

Polar bears have property X. (14)
Penguins have property X.

All animals have property X.

When given a forced choice between polar bears and antelopes versus polar bears and penguins, subjects judged the two animals from the same biological class, polar bears and antelopes, to be more similar than the two animals from different biological classes, polar bears and penguins. However, when asked to assess the inductive strength of each argument, argument (14) was judged to be less convincing than argument (13). That is, argument (13) had less diverse evidence, yet it was the stronger argument. Intuitively, although polar bears and penguins are from different biological classes, they still share the characteristic of living in a cold climate. It might seem than property X does not apply to all animals but only to animals living in cold climates.

Heit and Feeney (2005) investigated the non-diversity by property reinforcement effect further and came to a somewhat different conclusion. Essentially, their subjects judged polar bears and penguins to be more similar than polar bears and antelopes. Hence, when argument (13) was judged to be stronger than argument (14), Heit and Feeney's subjects were showing a diversity effect rather than a non-diversity effect. Heit and Feeney concluded that the diversity effect was indeed robust, and their results suggest that exceptions may be hard to show consistently.

The Developmental Approach

The first experiment on diversity-based reasoning in induction was actually a developmental one by Carey (1985), comparing six-year-olds and adults. Carey looked at patterns of inductive projection given the premises that two diverse animals, dogs and bees, have some biological property. The purpose of this study was to see whether subjects reason that "if two such disparate animals as dogs and bees" have this property then "all complex animals must" (p. 141). Indeed, adults made broad inferences to all animals, extending the property not only to things that were close to the premises (other mammals and insects) but also to other members of the *animal* category (such as birds and worms). In contrast, the children seemed to treat each premise separately; they drew inferences to close matches such as other mammals and insects, but they did not use the diversity information to draw a more general conclusion about animals. Therefore in this first attempt there was evidence for effects of diversity in adults but not children. In a follow-up study, Carey looked at diversity effects based on the concept *living thing* rather than *animal*. The key comparison was that children were taught a biological fact either about dogs and bees or about dogs and flowers, with the latter being even more diverse than the former. Given a fact about dogs and flowers, children did tend to generalize fairly broadly, suggesting that children may have some sensitivity to diversity of premise categories. However, if anything they tended to overgeneralize, extending the property not only to other living things but often to inanimate objects as well. Hence, there was suggestive evidence for the impact of diversity of premise categories in this study, although children did not show the same pattern as adults.

Continuing along this line of research that looks for diversity effects in children, Lopez et al. (1992) found limited evidence for nine-year-olds and no evidence for five-year-olds. For the five-year-olds, choices in a picture-based task did not show any sensitivity to diversity of premise categories, even when the diversity was explicitly mentioned by the experimenter. However, nine-year-olds did show sensitivity to diversity of premises, but only for arguments with a general conclusion category such as *animal* rather than a specific conclusion category such as *kangaroo*. Gutheil and Gelman (1997) attempted to find evidence of diversity-based reasoning for specific conclusions in nine-year-olds, using category members at lower taxonomic levels which would presumably enhance reasoning. However, like Lopez et al. (1992), Gutheil and Gelman did not find diversity effects in nine-year-olds, although in a control condition with adults, there was robust evidence for diversity effects.

More recently, however, Heit and Hahn (2001) reported diversity effects in children younger than nine years in experiments that used pictures of people and everyday objects as stimuli rather than animals with hidden properties. For example, children were shown a diverse set of dolls (a china doll, a stuffed doll, and a Cabbage Patch doll), all being played with by a girl named Jane. Also children were shown a non-diverse set, three pictures of Barbie dolls, being played with by Danielle. The critical test item was another kind of doll, a baby doll, and the question was who would like to play with this doll. In another stimulus condition, there was a diverse set of hats worn by one person, and a non-diverse set worn by another person, and again, the critical question was whether another hat would belong to the person with diverse hats or the person with non-diverse hats. For 74% of these critical test items, children age five to eight years made the diverse choice rather than the non-diverse choice. It seems from the Heit and Hahn experiments that children can follow the diversity principle at some level. However, it will take further work to establish the critical differences that led the past studies to not find diversity effects in children. (See Lo et al., 2002, for additional results, and Gelman, 2003, for further discussion.)

The Model-Based Approach

The psychological study of inductive reasoning has benefited from the development of computational models. Perhaps what is most impressive about these models is that they systematize what is inherently unsystematic, namely, probabilistic inference. Furthermore, these models capture important empirical phenomena in inductive reasoning, such as the diversity effect. In general, modeling work and experimental work on induction have been closely linked. In the following section, there are brief, conceptual descriptions of some of the earlier models and how they address the diversity effect. Several chapters in this book present newer modeling work in more detail (see Heit & Hayes, 2005, for another review of models of induction).

OSHERSON ET AL. (1990). The most influential model of category-based induction was developed by Osherson et al. (1990). This model has two main components. The first component assesses the similarity between the premise category or categories and the conclusion category. For example, argument (3) seems fairly strong because of the high level of similarity between dogs and wolves. In comparison, argument (15) below seems weaker, due to the lower similarity between dogs and parrots.

Dogs have hearts. (15)

Parrots have hearts.

The second component of the model measures how well the premise categories cover the superordinate category that includes all the categories mentioned in an argument. For example, in arguments (11) and (12), the relevant superordinate category is *mammal*. Intuitively, hippos and rhinos are somewhat similar mammals that do not span the whole superordinate category. In comparison, hippos and hamsters cover a broader range of mammals. The Osherson et al. model formalizes the calculation of coverage based on the structure of the *mammal* category. In effect, different instantiations of *mammal* are considered, from mice to cows to elephants, with the model looking for matches between the premise categories and these instantiations. Hippos and rhinos are a fairly good match to elephants and other large mammals, but they will be a poor match to smaller mammals. In comparison, the diverse set, hippos and hamsters, gives good matches to both large mammals such as elephants and small mammals such as mice. The diverse set is considered to have greater coverage because it yields a wider range of good matches with members of the superordinate category. More generally speaking, the similarity and coverage components of this model can be used to explain not only the diversity effect but also a dozen other phenomena in induction.

SLOMAN (1993). This model was implemented as a connectionist network, and perhaps the most important difference from the Osherson et al. (1990) model is that it does not have a separate component for assessing coverage of a superordinate category. Indeed, the Sloman model is valuable because it shows how much can be accomplished without this second mechanism, indeed addressing many of the same phenomena as the Osherson et al. model but without coverage. In brief, the way this model works is that premises of an argument are encoded by training the connectionist network to learn associations between input nodes representing the features of premise categories and an output node for the property to be considered. Then the model is tested by presenting the features of the conclusion category and measuring the activation of the same output node. For example, after the network has been trained that dogs have hearts, it will give a strong response if the features of *dog* are presented again as input. In addition, the network would give a strong output signal to wolves, indicating that wolves have hearts, because the featural representations of dogs and wolves would have a lot of overlap. The model accounts for diversity effects because training on a diverse set of categories will tend to strengthen a greater number of connections than training on a narrow range of categories. For example, training the network that both hippos and hamsters have a certain property would activate a broad range of features that apply to various mammals, leading to a strong conclusion that all mammals have that property. That is, hippos and hamsters would activate different features and different connections. In contrast, training the network that hippos and rhinos have a property would activate only a narrow range of features and connections. Although this model does have a notion of breadth of features, there is no distinct component for assessing coverage of a superordinate category, and indeed the model does not even rely on knowledge about superordinate categories. Nonetheless, the Sloman model can account for not only diversity effects but a variety of other phenomena.

HEIT (1998). The final model to be discussed is the Bayesian model by Heit (1998). This model is linked to eliminative accounts of hypothesis testing and as such is a normative model of how to reason with a hypothesis space. In addition, this account is fairly successful as a descriptive account in the sense that it predicts most of the same psychological phenomena as the Osherson et al. (1990) and Sloman (1993) models. According to the Bayesian model, evaluating an inductive argument is conceived of as learning about a property, in particular learning for which categories the property is true or false. For example, upon learning that dogs have some novel property X, the goal would be to infer which other animals have property X and which do not. For example, do wolves have the property and do parrots have the property? The key assumption is that for a novel property such as in this example, people would rely on prior knowledge about familiar properties to derive a set of hypotheses about what the novel property may be like. People already know a relatively large number of properties that are true of both dogs and wolves, so if property X applies to dogs, then it probably applies to wolves too. On the other hand, people know a relatively small number of properties that are true of both dogs and parrots. Hence property X is relatively unlikely to extend to parrots.

How does the Bayesian model explain diversity effects? In brief, diverse categories bring to mind a very different set of hypotheses than non-diverse categories. If hippos and hamsters have some novel property X in common, one considers familiar properties that are shared by all mammals, such as warm-bloodedness. Hence, property X too seems to extend broadly to all mammals, assuming that it is distributed in a similar way as other properties that are brought to mind. In contrast, if hippos and rhinos have property X in common, it is easier to think of familiar properties that are shared by hippos and rhinos but not most other mammals, such as being large and having a tough skin. Property X, too, may be distributed in the same way, namely, only to large, thick-skinned mammals, and seems less likely to extend to all mammals.

In sum, this section has illustrated different modeling approaches to induction, which have subsequently developed in later work. Interestingly, these three models address a similar set of phenomena with different theoretical frameworks, namely, in terms of hierarchically structured categories (Osherson et al., 1990), features in a connectionist network (Sloman, 1993), and beliefs within a hypothesis space (Heit, 1998). Not only is it possible to systematize what is unsystematic by nature, namely, probabilistic inference, but there is more than one way of doing so.

CONCLUSION

This chapter should at the very least illustrate the richness of research on induction. Research on this topic might seem to face a lot of challenges. After all, the degree of overlap with deduction has not yet been determined, and some accounts of reasoning simply define induction in terms of not being deduction. By their very nature, inductive inferences do not have a "right answer" in the same way as deductive inferences. Yet there are still regularities, such as the diversity principle, which can be studied from a variety of perspectives, including historical, philosophical, experimental, developmental, and computational. By no means is this regularity completely deterministic; indeed, there are well-documented exceptions to the diversity principle that are themselves of interest.

The material in this chapter should be seen as an invitation to consider different approaches to induction and different phenomena in induction, including those presented in the subsequent chapters of this book. All of the chapters refer to the experimental approach to at least some extent. The chapters by Hayes and by Medin and Waxman focus on the developmental approach. The chapters by Tenenbaum and Blok, Osherson, and Medin largely take the modeling approach. The chapters by Rips and Asmuth, and by Thagard, involve the historical approach, and the chapters by Oaksford and Hahn, and by Thagard, involve the philosophical approach. Finally, several of the chapters (by Rehder, Rips & Asmuth, Oaksford & Hahn, Thagard, and Feeney) directly address the question of what induction is.

References

Bacon, F. (1620–1898). *Novum organum.* London: George Bell and Sons.

Carey, S. (1985). *Conceptual change in childhood.* Cambridge, MA: Bradford Books.

Carnap, R. (1950). *Logical foundations of probability.* Chicago: University of Chicago Press.

Chater, N., & Oaksford, M. (2000). The rational analysis of mind and behavior. *Synthese, 122*, 93–131.

Earman, J. (1992). *Bayes or bust? A critical examination of Bayesian confirmation theory.* Cambridge, MA: MIT Press.

Evans, J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning.* Hove, UK: Psychology Press.

Feeney, A., & Handley, S. J. (2000). The suppression of q-card selections: Evidence for deductive inference in Wason's selection task. *Quarterly Journal of Experimental Psychology, 53A*, 1224–1242.

Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought.* New York: Oxford University Press.

Goel, V., Gold, B., Kapur, S., & Houle, S. (1997). The seats of reason: A localization study of deductive and inductive reasoning using PET (O15) blood flow technique. *NeuroReport, 8*, 1305–1310.

Goodman, N. (1972). *Problems and projects.* Indianapolis: Bobbs-Merrill.

Gorman, M. E., Tweney, R. D., Gooding, D. C., & Kincannon, A. P. (Eds.), (2005). *Scientific and technological thinking.* Mahwah, NJ: Lawrence Erlbaum Associates.

Gutheil, G., & Gelman, S. A. (1997). Children's use of sample size and diversity information within basic-level categories. *Journal of Experimental Child Psychology, 64*, 159–174.

Harman, G. (1999). *Reasoning, meaning, and mind.* Oxford: Oxford University Press.

Heit, E. (1998). A Bayesian analysis of some forms of inductive reasoning. In M. Oaksford & N. Chater (Eds.), *Rational models of cognition*, 248–274. Oxford: Oxford University Press.

Heit, E. (2000). Properties of inductive reasoning. *Psychonomic Bulletin & Review*, 7, 569–592.

Heit, E., & Feeney, A. (2005). Relations between premise similarity and inductive strength. *Psychonomic Bulletin & Review, 12,* 340–344.

Heit, E., & Hahn, U. (2001). Diversity-based reasoning in children. *Cognitive Psychology, 43*, 243–273.

Heit, E., Hahn, U., & Feeney, A. (2005). Defending diversity. In W. Ahn, R. Goldstone, B. Love, A. Markman, & P. Wolff (Eds.), *Categorization inside and outside of the laboratory: Essays in honor of Douglas L. Medin*, 87–99. Washington, DC: American Psychological Association.

Heit, E., & Hayes, B. K. (2005). Relations among categorization, induction, recognition, and similarity. *Journal of Experimental Psychology: General, 134*, 596–605.

Heit, E., & Rotello, C. M. (2005). Are there two kinds of reasoning? In *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Heit, E., & Rubinstein, J. (1994). Similarity and property effects in inductive reasoning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 411–422.

Hempel, C. G. (1966). *Philosophy of natural science.* Englewood Cliffs, NJ: Prentice Hall.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: MIT Press.

Horwich, P. (1982). *Probability and evidence.* Cambridge, UK: Cambridge University Press.

Howson, C., & Urbach, P. (1993). *Scientific reasoning: The Bayesian approach.* Chicago: Open Court.

Hume, D. (1777). *An enquiry concerning human understanding.* Oxford: Clarendon Press.

Johnson-Laird, P. (1983). *Mental models.* Cambridge, MA: Harvard University Press.

Johnson-Laird, P. N. (1994). Mental models and probabilistic thinking. *Cognition, 50*, 189–209.

Johnson-Laird, P. N., Legrenzi, P., Girotto, V., Legrenzi, M. A., & Caverni, J. P. (1999). Naive probability: A mental model theory of extensional reasoning. *Psychological Review, 106*, 62–88.

Lo, Y., Sides, A., Rozelle, J., & Osherson, D. (2002). Evidential diversity and premise probability in young children's inductive judgment. *Cognitive Science, 16*, 181–206.

Lopez, A., Gelman, S. A., Gutheil, G., & Smith, E. E. (1992). The development of category-based induction. *Child Development, 63,* 1070–1090.

Lopez, A., Atran, S., Coley, J. D., Medin, D. L., & Smith, E. E. (1997). The tree of life: Universal and cultural features of folkbiological taxonomies and inductions. *Cognitive Psychology, 32,* 251–295.

Medin, D. L., Coley, J. D., Storms, G., & Hayes, B. K. (2003). A relevance theory of induction. *Psychonomic Bulletin & Review, 10*, 517–532.

Nagel, E. (1939). *Principles of the theory of probability*. Chicago: University of Chicago Press.

Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review, 101*, 608–631.

Osherson, D. N., Smith, E. E., Wilkie, O., Lopez, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97*, 185–200.

Osherson, D. N., Perani, D., Cappa, S., Schnur, T., Grassi, F., & Fazio, F. (1998). Distinct brain loci in deductive versus probabilistic reasoning. *Neuropsychologia, 36*, 369–376.

Poletiek, F. (2001). *Hypothesis testing behaviour*. Hove, UK: Psychology Press.

Popper, K. R. (1963). *Conjectures and refutations: The growth of scientific knowledge.* London: Routledge.

Proffitt, J. B., Coley, J. D., & Medin, D. L. (2000). Expertise and category-based induction. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 811–828.

Rips, L. J. (2001). Two kinds of reasoning. *Psychological Science, 12*, 129–134.

Salmon, W. C. (1984). *Scientific explanation and the causal structure of the world.* Princeton, NJ: Princeton University Press.

Shepard, R. N. (1987). Towards a universal law of generalization for psychological science. *Science, 237*, 1317–1323.

Skyrms, B. (2000). *Choice and chance: An introduction to inductive logic.* (Fourth edition). Belmont, CA: Wadsworth.

Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology, 25*, 231–280.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*, 3–22.

Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning.* Mahwah, NJ: Lawrence Erlbaum Associates.

Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*, 629–641.

Wayne, A. (1995). Bayesianism and diverse evidence. *Philosophy of Science, 62*, 111–121.

Weisberg, R. W. (1986). *Creativity: Genius and other myths.* New York: Freeman.