

## 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.



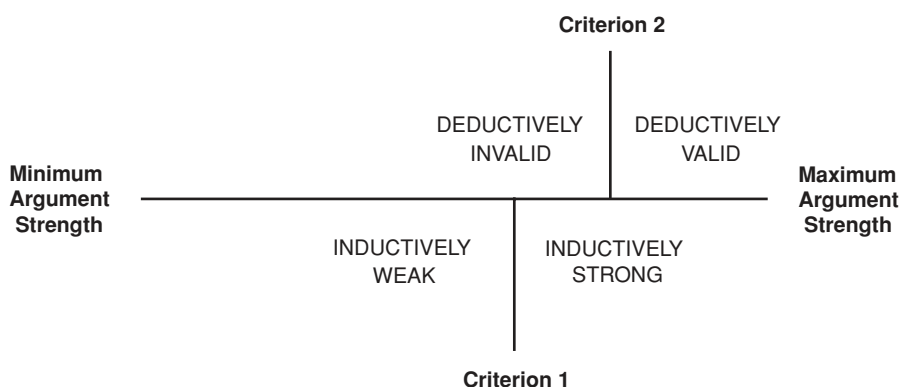


FIGURE 1.1. Criterion-shift account of deduction and induction.

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