

Evidence and Inductive Inference

Nevin Climenhaga

Forthcoming in *The Routledge Handbook of the Philosophy of Evidence*, ed. Maria Lasonen-Aarnio and Clayton Littlejohn (Routledge, 2020)

Abstract

This chapter presents a typology of the different kinds of inductive inferences we can draw from our evidence, based on the explanatory relationship between evidence and conclusion. Drawing on the literature on graphical models of explanation, I divide inductive inferences into (a) downwards inferences, which proceed from cause to effect, (b) upwards inferences, which proceed from effect to cause, and (c) sideways inferences, which proceed first from effect to cause and then from that cause to an additional effect. I further distinguish between direct and indirect forms of downwards and upwards inferences. I then show how we can subsume canonical forms of inductive inference mentioned in the literature, such as inference to the best explanation, enumerative induction, and analogical inference, under this typology. Along the way, I explore connections with probability and confirmation, epistemic defeat, the relation between abduction and enumerative induction, the compatibility of IBE and Bayesianism, and theories of epistemic justification.

1. Inductive Inference

There are two senses of ‘evidence.’ The first is the “having evidence” sense. When we *have* A as evidence, A is part of a body of evidence we possess. The second sense is that of “evidence for.” When A is evidence *for* B, A *confirms* B—that is, raises its probability. When we have A as evidence and A is evidence for B, we are in a position to *infer* from A to B—concluding that B is true, or that B is probable, or that B is more probable than it would otherwise be.

The inferences we draw from our evidence are of different kinds. Traditionally, inferences are divided into *deductive* and *inductive*. In a valid deductive argument, it is impossible for the premises to be true and the conclusion false; so if we are sure of A and deductively infer B from A, we can be sure of B as well.¹ While philosophers have focused much of their energy on understanding deductive inference, our everyday inferences are more

¹ At least, this is plausibly true in ideal cases of competent deductive inference. Schechter (2013) discusses difficulties with formulating a general principle here.

often inductive in form. In a good inductive argument, our premises support our conclusion to some degree, but it is possible for the premises to be true and the conclusion false.

Consequently, inductive inferences only let us move from evidence to conclusion with some degree of probability.

This last point requires clarification, as the literature is divided on exactly how a good inductive inference supports its conclusion. On one conception, an inductive inference should make its conclusion more probable than not (see, e.g., Hurley 2006: 44-45). On another conception, it need only make its conclusion more probable than it would otherwise be (see, e.g., Carnap 1950: 205-06). The former fits with a conception of inference more generally as a cognitive process that results in belief.² The latter fits with a conception of inference more generally as a cognitive process that includes not only changes in first-order beliefs, but also changes in degrees of belief and/or beliefs about probabilities.³ It also makes inductive *inferences* more closely analogous to inductive *arguments*, inasmuch as the latter category includes arguments that confirm, but do not make more probable than not, their conclusions (Swinburne 2004: 4-6).⁴

This chapter focuses on the question of what kinds of inductive inferences there are, with the goal of helping us better understand and improve our ordinary inferential practices. In order to include more reasoning under the scope of this discussion, I adopt the second conception of inductive inference above, on which it includes any reasoning from premises to

² For example, according to Sanford (2015: 509), “Inference occurs only if someone, owing to believing the premises, begins to believe the conclusion or continues to believe the conclusion with greater confidence than before.”

³ Bayesians frequently use the term ‘inference’ in this way. For example, Levi (1977: 6-7) characterizes the outcome of a “direct inference” (see Section 3.1.1 below) as “assign[ing] a degree of personal or credal probability to the hypothesis that [the event] *e* results in an event of kind *R*.”

⁴ Logic textbooks differ on whether to measure inductive argument strength by the probability of the conclusion conditional on the premises, or the degree to which the premises confirm the conclusion. Some empirical research suggests that folk evaluations of the strength of inductive arguments better tracks the latter (see Crupi et al. 2008 for discussion and references).

a conclusion (taken to be) confirmed by that conclusion.

Little formulaic work in the past century has aimed at giving a typology of different kinds of inductive inference. But the following are among the most common inductive inference forms mentioned in introductory logic and critical reasoning textbooks, encyclopedia articles, and philosophical discussions:⁵

- Abduction
- Analogical inference
- Bayesian inference
- Causal inference
- Direct inference
- Enumerative induction
- Inverse inference
- Inference to the best explanation
- Predictive inference
- Statistical inference
- Universal inference

This chapter proceeds as follows. First, in Section 2, I draw on the technical literature on graphical models of explanation to give a principled typology of inductive inferences. Then, in Section 3, I classify the above forms within this typology. Finally, in Section 4, I note some philosophical implications of this typology.

2. A Typology of Inductive Inference

We are wondering whether various members of an extended family are smokers. We know that the only causal factor that influences whether someone smokes is parental habits: if a parent smokes, their child is more likely to smoke than if that parent does not smoke. We have the partial family tree given in Figure 1, with the relatives denoted by their relation to one individual, Jane. Suppose we learn that Father smokes. What can we infer?

⁵ For introductory textbooks, I consulted Gensler 2002, Hurley 2006, Vaughn 2009, and Copi et al. 2014. For encyclopedias, I consulted Flew 1979, Audi 2015, and the *Stanford Encyclopedia of Philosophy*. As for philosophical discussions, I rely especially on Carnap 1950 below, and cite others when relevant.




Figure 1

First, we might infer that Jane smokes. We cannot infer this with certainty. It is possible that Father smokes but has not passed this habit on to Jane. But that Father smokes makes it more likely that Jane smokes.

Second, we might infer that Grandfather smokes. We cannot infer this with certainty. It is possible that Father picked up smoking without learning it from Grandfather. But Grandfather's smoking is one way we might explain Father's smoking, and that Father smokes thus makes it more likely that Grandfather smokes.

Third, we might draw further inferences from our two preliminary conclusions. Just as Father's smoking makes it more likely that Jane smokes, Grandfather's smoking makes it more likely that Uncle smokes, and Jane's smoking makes it more likely that Daughter smokes. So we can further infer (with some probability) that Daughter smokes and that Uncle smokes.

There is one further inference we cannot make. We cannot infer from the preliminary conclusion that Jane smokes that Mother smokes. This is because our only reason to think that Jane smokes is that Father's smoking might cause Jane to smoke. We cannot say, "Father

smokes, which will probably lead to Jane smoking, which is probably explained by Mother's smoking." Since the only reason we have to think that Jane smokes is that Father's smoking predicts this, there is nothing residual to account for that Mother's smoking could help explain.

This example lets us divide inductive inferences into two kinds: direct and indirect. Our first two inferences were direct: we moved from Father's smoking to Jane's smoking and Grandfather's smoking without any intermediate steps. Our last two (legitimate) inferences were indirect: we moved from Father's smoking to Daughter's smoking and Uncle's smoking, but only by moving through Jane's and Grandfather's smoking.

Within direct inferences, we can distinguish *upwards inferences* from *downwards inferences*. Downwards inferences move from cause to effect—parent to child—and upwards inferences move from effect to cause—child to parent.⁶

Within indirect inferences, we can distinguish different combinations of direct inferences, with the direct inferences starting either at our evidence or at the conclusion of an earlier direct inference. Combinations of two or more upwards inferences are *indirect upwards inferences*. Combinations of two or more downwards inferences are *indirect downwards inferences*. Finally, if we begin with a (direct or indirect) upwards inference followed by a (direct or indirect) downwards inference, we have a *sideways inference*. Sideways inferences let us move from an observed effect to other hypothetical effects of an underlying explanation.

This lets us construct the typology in Table 1.⁷ We saw above that we cannot first infer downwards from a cause to an effect, and then from that effect upwards to another

⁶ For exposition, I use the language of "cause" and "effect," but I mean this to generalize to explanatorily priority relations more generally, as I explain below.

⁷ Swinburne (2001: ch. 4) uses the terms 'downward inference' and 'upward inference' in a parallel sense. The term 'sideways inference' is my own, but Swinburne describes (without naming) this kind of inference in a similar way.

cause. This means we cannot add an upwards inference at the end of either a downwards inference or sideways inference to get a new kind of indirect inference, and our typology is thus complete.

	<i>Inference Type</i>	<i>Example</i>
<i>Direct</i>	Downwards	Father→Jane
	Upwards	Father→Grandfather
<i>Indirect</i>	Downwards (downwards + downwards)	Father→Jane→Daughter
	Upwards (upwards + upwards)	Daughter→Jane→Father→Grandfather
	Sideways (upwards + downwards)	Jane→Father→Grandfather→Uncle

Table 1

We arrived at this typology through consideration of a special case. To generalize it, we can adopt the following framework for thinking about inferences:⁸

- (1) Different facets of the world which we are interested in reasoning about can be organized into *variables*.

Variables correspond to questions we can ask about the world: “Does Jane smoke?” “How often does Jane exercise?” “What is Jane’s blood pressure?” Answers to these questions (“Yes,” “Twice a week,” “120 over 80”) correspond to *values* these variables can take on.

- (2) We can model variables as the fundamental relata of inference.

We can infer from the observation that Jane’s blood pressure is high to the conclusion that she does not exercise regularly. But we can also infer more generally from Jane’s blood pressure to her exercise habits. We can infer from one variable to another when the two are *probabilistically dependent*: learning the value of one changes the probabilities of different values of the latter. (This builds on the broad conception of inference adopted in Section 1, on which inference includes changes in credences/beliefs about probabilities as well as

⁸ See Climenhaga 2020 for a fuller elaboration of this framework.

changes in first-order beliefs.)

- (3) These variables can be organized into a directed acyclic graph (DAG) that represents the explanatory relations between the variables and obeys the Markov condition.

A DAG represents explanatory relations between our variables by use of directed arrows. We include an arrow from one variable X to another variable Y iff what value X takes on directly influences what value Y takes on. This influence could be causal, as in Figures 1 and 2. But arrows can also represent non-causal influence, as when the value X takes on grounds or partially grounds the value Y takes on.⁹ Downwards, upwards, and sideways inferences can accordingly track non-causal explanatory relations as well as causal relations. In Figure 3, Socrates' parents' existence is causally prior to Socrates' existence, and Socrates' existence metaphysically prior to {Socrates, Plato}'s existence—that is, whether Socrates exists helps metaphysically determine whether the set {Socrates, Plato} exists. This lets us infer directly downwards both from the existence of Socrates' parents to the existence of Socrates, and from the existence of Socrates to the existence of {Socrates, Plato}.




Figure 2




Figure 3

X is a *parent* of Y in a DAG iff there is an arrow from X to Y , and an *ancestor* of Y iff it is a parent, or parent of a parent, etc. If X is a parent/ancestor of Y , Y is a child/descendant of X . A DAG obeys the *Markov condition* just in case a variable's parents screen it off from all non-descendants. Formally:

A DAG obeys the *Markov condition* iff for any variable X , X is probabilistically

⁹ In the past, DAGs have primarily been used to model causal priority (e.g., Pearl 2000, Spirtes et al. 2000). Schaffer (2016) defends the use of DAGs to represent metaphysical priority. See Climenhaga 2020: sec. 2.3 for further discussion and references.

independent, given its parents, from any other conjunction of non-descendants Z. Informally, the Markov condition says that if Y already tells us everything relevant to predicting X in advance, the only way to get further evidence about the value of X is by learning about its effects. For example, if we know that the only thing that directly causally influences whether Jane smokes is whether her parents smoke, then if we know that Father smokes and Mother doesn't, learning that Grandfather smokes tells us nothing. The only evidence we can get about her smoking habits are possible effects of those habits—like her children smoking or her blood pressure.

If our DAG obeys the Markov condition, then the probability of Y given X is only different from the unconditional probability of Y if we can infer from X to Y in one of our three ways: upwards, downwards, or upwards + downwards.¹⁰ A metaphor: if we imagine the nodes of a DAG as individuals in a family tree (in a species without sex and with no restriction on the number of parents or the possibility of incestuous reproduction), our typology of inference says that evidence about one node gives us information about other nodes that are *related to that node by blood, but not those only related by marriage*. Just as I can draw inferences from the results of your genetic test about your blood relations, but not your in-laws, I can draw inferences from observing the value of some variable only about the “blood relations” of that variable.

This result holds only in situations in which we have no background evidence (that is, our evidence does not include the values of any of the other variables in our DAG). Fully exploring the impact of background evidence on inferences is beyond the scope of this

¹⁰ This follows from Theorem 1.2.4 and the definition of directional separation in Pearl 2000: 16-18. There are several further connections between probability and our typology. In Climenhaga 2020 and ms-b I argue that which theorem of probability we should use to calculate the probability of A given B depends on the explanatory relation between A and B. Another question worth further exploration is whether, on standard probabilistic measures of confirmation, confirmation lessens as inferences become more indirect. For example, it is plausible that in Figure 1, the proposition that Jane smokes must confirm that Father smokes more strongly than it confirms that Grandfather smokes.

chapter, but an example will suffice to show the complications that can arise.

Consider Figure 2. Suppose we observe that Jane has high blood pressure, and infer that she does not exercise regularly. We subsequently learn that Jane is a smoker. Relative to the background observation that Jane has high blood pressure, Jane’s smoking raises the probability that Jane exercises regularly—even though neither of these factors influences the other, and they have no common causes. Because our new evidence “explains away” the initial observation that Jane has high blood pressure, there is less explanatory work for lack of exercise to do; this evidence thus undermines the support that high blood pressure gives to lack of exercise. Philosophers have referred to this kind of undermining as “undercutting defeat” (Pollock 1986) and “explaining away” (Schupbach 2016). It is worth noting, though, that in other cases antecedent knowledge of factors that influence an observation can strengthen the inference from it to another explanation—and here talk of “undercutting” or “explaining away” is not apt. For example, if we had learned that Jane does *not* smoke, this would have strengthened the inference from high blood pressure to lack of exercise, by removing a possible alternative explanation of this observation.

3. Classifying Inductive Inference Forms Under this Typology

In this section I return to the assortment of inference kinds compiled in the introduction, and bring them under the typology developed above.

3.1 Downwards inference

The first kind of inference in our typology is downwards inference—from cause to effect, or explanation to prediction.

3.1.1 Direct inference

The term ‘direct inference’ has been used to describe inferences from statistical hypotheses to specific events (Carnap 1950: 207, Henderson 2020)—that is, inferences from

a population to a sample.¹¹ An inference from a given proportion of white balls in an urn to the probability of drawing some number of white balls in some number of samples is a direct inference in this sense. The term has also been used to describe inferences from physical chances to outcomes (Levi 1977)—for example, an inference from the hypothesis that a coin is fair to a prediction about the number of heads that will be flipped in a series of trials.

Hypotheses about physical chances are prior to particular outcomes, and hypotheses about the characteristics of a population are prior to observations of individuals sampled from that population. (That is, statistical generalizations are prior to our *observations* of individuals. They are not usually prior to the features of the individuals themselves. See Section 3.2.1.1 below.) Direct inferences in both these senses are thus downwards inferences. These are typically “direct” in the sense I have been using the term as well, in that the inferences are about children of the chance/statistical hypothesis in the explanatory network, so that they are not mediated by any intermediary factors.

3.1.2 Other kinds of downwards inference

Not all direct (in our sense) downwards inferences proceed from a statistical description of a population or a chance hypothesis about a process to a statement about some member of that population or outcome of that process. When you infer from your friend’s having promised to meet you for lunch at noon that they will be at the restaurant at noon, you are inferring from a cause to an effect, but the cause is neither a statistical distribution nor a hypothesis ascribing precise chances to a physical process.

We could subdivide downwards inference into inference from *quantitative* hypotheses (e.g., about statistical distributions or chance physical processes) and inference from *qualitative* hypotheses (e.g., about the actions of rational agents). Inferences of the former kind are what others have called “direct inference.” They are of special interest because they

¹¹ The term ‘statistical syllogism’ is sometimes used here as well (e.g., Gensler 2002: ch. 13).

apparently allow for the straightforward derivation of precise conditional probabilities. But inferences from qualitative hypotheses are also common, in both science and ordinary life.

3.2 Upwards inference

The second kind of inference is upwards inference—from effect to cause, or observation to explanation.

3.2.1 Inverse inference/Bayesian inference/statistical inference

“Direct inference” in the sense of Section 3.1.1 was traditionally contrasted with “inverse inference”—inference from a sample to a population, or from effects to causes. Bayes’ Theorem was originally formulated to deal with problems of inverse inference (see Fienberg 2006). If we take the older ‘inverse inference’ and the more recent ‘Bayesian inference’ to refer to all inferences that use Bayes’ Theorem, then these are arguably equivalent to upwards inference in the sense of this chapter (see Climenhaga 2020 and ms-b).

Some authors use the term ‘inverse inference’ only for narrower kinds of inference, such as “inference from a sample to a population” (Carnap 1950: 207). Such inferences are also called “statistical inferences”—i.e., inferences to a statistical hypothesis about a population (Howson and Urbach 2006).

One natural way to subdivide upwards inference, corresponding to the distinction between different kinds of downwards inference above, is with respect to whether their conclusions are qualitative or quantitative. Inferences to statistical or chance hypotheses are the inverse of “direct inferences” in the sense of Section 3.1.1, while inferences to qualitative hypotheses are the inverse of other kinds of downwards inferences.

3.2.1.1 Enumerative induction/universal inference

A special case of statistical inference that has received disproportionate philosophical attention is what Carnap (1950: 208) calls “universal inference.” In universal inference, one infers from the fact that some sample of a population has a feature to the conclusion that the

whole population has that feature. For example, one infers from the fact that all observed ravens have been black to the conclusion that all ravens are black.

Discussion of universal inference goes all the way back to Aristotle, and it was historically taken as the paradigm of induction more generally (see, e.g., Flew 1979: 171). It was classically called “induction by simple enumeration,” or “enumerative induction” (see Norton 2010 for historical references). Over time the meaning of these terms has shifted, though, and today they are often used in a broader sense that covers predictive inference (see Section 3.3.1) as well (e.g., Russell 1948, Henderson 2020).

On the present typology, universal inference is a special kind of upwards inference to a quantitative hypothesis: one where the quantitative hypothesis is that 100% of the population has the feature in question. In classifying universal inference as a form of upwards inference, I side with Harman (1965), who held that enumerative induction is a special case of inference to the best explanation (see Section 3.2.2 below). This thesis has proven controversial, with the controversy largely turning on the question of whether universal statements (like “all ravens are black”) really *explain* their instances (like “this raven is black”) (Ennis 1968). On the plausible view that generalizations are grounded in their instances, the direction of explanation goes the other way: all ravens are black *because* this raven is black, that raven is black, and so on.¹² However, as Weintraub (2013: 211-12) observes, even if generalizations do not explain their instances, they may explain our *observations*. In canonical enumerative inductions, the evidence is about an individual or group *sampled* from the population. Here the evidence is not that this raven is black, that that raven is black, etc., but that we have observed this black raven, that we have observed that black raven, etc. And while the fact that all ravens are black may not explain why particular

¹² One response to this is that *lawlike* generalizations explain their instances (Harman 1968: 531, Lipton 2004: 97, Bhogal 2017). As Weintraub (2013: 211-12) argues, though, not all universal inferences are to laws.

ravens are black, it does explain why we have observed only black ravens.

The situation is different if we start with a population whose members we are already acquainted with, and then infer from one of them having a feature to others sharing that feature. For example, call Jane, Mother, and Father, “Jane’s immediate family.” We already know everyone in Jane’s family, and then we learn that Jane is a smoker. From this we conclude that everyone in Jane’s immediate family smokes. While this inference fits the logical form presented in textbooks, it is not the kind of inference philosophers usually have in mind when discussing enumerative induction. It does not proceed directly upwards from a sample to a population, but upwards from Jane to Mother and Father, and then downwards from these instances to the generalization that all these people smoke. It is accordingly a kind of sideways inference.

3.2.2 Abduction/inference to the best explanation

Another commonly discussed form of inductive inference is *abduction*.¹³ This term is variously used (see Douven 2017). Some authors use ‘abduction’ to describe the *invention* of a hypothesis, contrasting this with the selection of a hypothesis. I focus here solely on abduction in the inferential sense. Harman (1965) equates abduction in this sense with “inference to the best explanation” (IBE). Much discussion of IBE has focused on difficulties relating to inferring to the *best* explanation (Thagard 1978, van Fraassen 1989, Lipton 1993, Climenhaga 2017b). If we set these complications aside, and use ‘abduction’ and ‘IBE’ to describe any inference from a proposition to a possible explanation of that proposition, then it is identical to upwards inference in our sense.

With that said, the examples used to illustrate IBE in the literature almost invariably involve qualitative hypotheses. Consider this stock of examples from Lipton’s (2004: 56) celebrated book:

¹³ This term was coined by C.S. Peirce, who also used the term ‘retroduction’ for this form of inference.

The sleuth infers that the butler did it, since this is the best explanation of the evidence before him. The doctor infers that his patient has measles, since this is the best explanation of the symptoms. The astronomer infers the existence and motion of Neptune, since that is the best explanation of the observed perturbations of Uranus. Chomsky infers that our language faculty has a particular structure because this provides the best explanation of the way we learn to speak. Kuhn infers that normal science is governed by exemplars, since they provide the best explanation for the observed dynamics of research.

None of these five inferences are to statistical or chance hypotheses, although one of them (the existence of Neptune) explains the observed data (the perturbations of Uranus) by means of a mathematical derivation from this hypothesis (together with Newton's laws and background evidence about the location of the other planets).

That standard examples used to illustrate abduction are upwards inferences to qualitative hypotheses suggests that the long-running debate about the compatibility of IBE and Bayesianism (Henderson 2014) is closely related to more general questions about the relation of quantitative and qualitative reasoning. In particular, Bayesian inferences will be abductive inferences only if qualitative explanatory considerations are relevant to probabilities (contra van Fraassen 1989), and abductive inferences will be Bayesian inferences only if probabilities can legitimately be assigned to qualitative hypotheses (contra Haack 2018).

3.2.3 Causal inference

“Causal inference” is frequently described as a special form of inference relevant to such areas as the social and medical sciences, where we are given several variables but do not know—and thus must infer—the causal relationships among them (Spirtes et al. 2000, Climenhaga et al. forthcoming). Introductory logic textbooks that discuss inductive reasoning frequently devote a section to “causal reasoning,” usually focusing on Mill's (1843) methods for identifying, in a set of data, causal relationships between different factors (see, e.g., Copi et al. 2014: ch. 12.4, Gensler 2002: ch. 13.7, Vaughn 2009: ch. 8, Hurley 2006: ch. 9.2).

So far this chapter has discussed cases where we are given a causal network, and

make inferences relative to that network. Causal inference is needed in cases where we are uncertain what the explanatory relationships between different variables are. These two situations are opposite sides of the same coin. Suppose you are uncertain whether smoking would make Jane more likely to develop cancer. Call the causal network on which it would the dependent network and the one on which it would not the independent network. Now suppose you observe that Jane smokes and has cancer. This observation is more likely given the dependent network than given the independent network. Hence, this observation is evidence for the dependent network; hence, you can infer (with some probability) from this observation to the dependent network (cf. Climenhaga 2017a).

Since causal inferences are inferences from predictions to explanations, causal inferences are a kind of upwards inference. The causal relationship between smoking and cancer is prior to both variables, in that what difference smoking makes to cancer itself makes a difference to how likely it is that someone both smokes and develops cancer. In observing both variables and drawing an inference about the causal network, one is inferring upwards to an explanation of one's evidence.

Hypotheses about explanatory relationships may be quantitative (assigning precise conditional probabilities, as in stock balls-and-urns models used to illustrate Bayesian inference), or qualitative (specifying relations of explanatory priority but not precise conditional probabilities); so causal inferences may be inferences to either quantitative or qualitative hypotheses.

3.3 Sideways inference

The third kind of inference is sideways inference: inference upwards from effect to cause or explanation, and then downwards from that cause or explanation to another effect.

3.3.1 Predictive inference

Carnap (1950: 207) defines a predictive inference as “an inference from one sample to

another sample not overlapping with the first.” The evidence is about observed members of some population, and the conclusion about unobserved members of that population—e.g., the next member one observes. Hume (1748/1999) presents canonical examples of this inference form:

- The bread I have formerly eaten nourished me; therefore, the next bread I eat will nourish me.
- All the snow I have felt has been cold; therefore, that snow I see falling from the sky will be cold.

Hume saw predictive inference as the central case of induction, and subsequent philosophical theorizing has placed it alongside universal inference as one of the paradigmatic forms of induction.¹⁴ Carnap calls predictive inference “the most important and fundamental kind of inductive inference” (1950: 207), and others who have sought a universal “inductive rule” have tended to follow him on this (see Climenhaga ms-a).

On the typology advanced here, predictive inference is just one kind of indirect sideways inference—an inference that proceeds upwards through an explanatory hypothesis and then downwards from that to a prediction of that hypothesis. Like universal or statistical inference, we can understand predictive inference as proceeding from information about members of a group we have *sampled* or *observed*. We move from the bread we have sampled or the snow we have felt upwards to a generalization about most bread nourishing, or all snow being cold, and then downwards to the prediction that the next bread eaten will be nourishing, or the next snow felt cold. Standard discussions of predictive inference, and especially attempts to give rules for quantifying the strength of predictive inference based solely on factors like number of samples, are best explicable under this sampling paradigm.

But we can also infer from observed instances to unobserved instances in other ways. We could infer from the bread we have eaten nourishing to an explanation in terms of what

¹⁴ See, for example, the vast literature on Goodman’s (1954) “grue” problem.

Hume called bread's "secret powers", and then downwards from this to the prediction that other bread nourishes as well. This looks more like an argument from analogy: we reason that unobserved bread that resembles observed bread in "colour, weight and consistency" (Hume 1748/1999: sec. 4.2) will also resemble it in being nourishing.

3.3.2 Analogical inference

In an analogical inference, one infers from the fact that a is both F and G , and that b is F , that b is also G (where F and G may be atomic or conjunctive predicates). Here are two examples from the *Stanford Encyclopedia of Philosophy* entry on analogical reasoning:

Reid notes a number of similarities between Earth and the other planets in our solar system: all orbit and are illuminated by the sun; several have moons; all revolve on an axis. In consequence, he concludes, it is "not unreasonable to think, that those planets may, like our earth, be the habitation of various orders of living creatures" (1785: 24) (Bartha 2019: sec. 1).

In 1934, the pharmacologist Schaumann was testing synthetic compounds for their anti-spasmodic effect. These drugs had a chemical structure similar to morphine. He observed that one of the compounds—*meperidine*, also known as *Demerol*—had a physical effect on mice that was previously observed only with morphine: it induced an S-shaped tail curvature. By analogy, he conjectured that the drug might also share morphine's narcotic effects (Bartha 2019: sec. 2.1).

In these examples, we first infer upwards from Fa and Ga to some causal relationship between Fa and Ga . This is a *causal inference* in the sense of Section 3.2.3: inference to a causal network. In Reid's argument for life on other planets, we infer from Earth's (a) orbiting the sun, (b) having a satellite, (c) revolving on an axis, and (d) having life that (a)-(c) causally contribute to (d) on Earth. In Schaumann's conjecture that meperidine is a pain-killer, we infer upwards from the presence of tail curvature and pain-killing among the effects of morphine that these two phenomena have a common cause in morphine's chemical structure. Second, we infer from this causal network for a to an analogous network for b —to (a)-(c) similarly influencing whether (say) Mars has life; and to meperidine's chemical structure similarly influencing whether it has tail-curving and pain-reducing effects. This inference is itself a sideways one, proceeding through some (perhaps implicit) higher-order

hypothesis about a common explanation of these two networks—for example, underlying laws of nature. Finally, relative to the background evidence that Fb , we infer downwards from this second network to Gb —observing (a)-(c) in Mars, we conclude that it will also have life; observing the tail curvature effect of meperidine, we conclude that it has a similar chemical structure and therefore will also be an effective pain-killer.

As Bartha's (2019) other examples illustrate, a wide variety of arguments are classified as analogical in the literature. I lack the space to consider further examples here, but I suspect that we can understand most analogical inferences as sideways inferences of the above form, comprising two upwards steps followed by two downwards steps. We do need to generalize this form in one way, however. In the examples above, the first upwards inference is to a causal network relating Fa and Ga (more precisely: a causal network relating variables taking Fa and Ga as values), and the first downwards inference is to a similar network for b . Bartha (2019: sec. 3.3.2) notes (against Hesse 1966) that not all good analogical arguments posit a causal relationship between F and G —for example, analogical arguments in mathematics do not. We can account for this by also allowing for inferences to non-causal explanatory networks—e.g., a network on which whether a is F grounds whether a is G . This allows for justified analogical inferences not only in empirical reasoning but also in such domains as mathematics and ethics.

3.3.3 Other kinds of sideways inference

While predictive and analogical inference are the varieties of sideways inference most widely discussed, the category is wider than just these two. Consider an inference from South America and Africa having complementary shapes to the theory of continental drift, and from there to the probable existence of undersea rifts where plates meet. This inference relies neither on an analogy between shapes of continents and undersea rifts nor on these two things being part of some population we are sampling from. But it is, like analogical and predictive

inferences, a sideways inference—proceeding upwards from our evidence to an explanatory hypothesis, and downwards from there to a prediction of that hypothesis.

4. Conclusion

In Section 2, I defended a typology of inductive inference based on the explanatory relationship between premises and conclusion. There I argued that inferences can proceed either downwards, from cause to (direct or indirect) effect, upwards, from effect to (direct or indirect) cause, or sideways, from effect to cause to additional effect. In Section 3, I classified canonical forms of inductive inference under this typology. Table 2 summarizes these results.

<i>Inference Type</i>	<i>Canonical Inductive Inference Forms</i>
Downwards	<ul style="list-style-type: none"> • Direct inference
Upwards	<ul style="list-style-type: none"> • Inverse inference/Bayesian inference/statistical inference <ul style="list-style-type: none"> ◦ Special case: universal inference/enumerative induction • Abduction/inference to the best explanation • Causal inference
Sideways	<ul style="list-style-type: none"> • Predictive inference • Analogical inference

Table 2

The typology defended in this chapter is not philosophically neutral. I have mentioned connections above with probability theory, confirmation theory, epistemic defeat, the relation between abduction and enumerative induction, and the compatibility of IBE and Bayesianism. In closing I note one final implication.

The model of inference presented here fits most naturally with a form of foundationalism about epistemic justification, on which your evidence is fixed prior to your inferences and determines what inferences you can rationally draw. Recall that our typology precludes inferring downwards from evidence A to B and then upwards from B to C. You cannot infer from Father’s smoking to Jane’s smoking, and then from Jane’s smoking to Mother’s smoking. But if your evidence was B rather than A, you could infer both A and C if

these both predict B; you could infer to both Mother and Father smoking from Jane's smoking.

This characteristic of rational inference is difficult to reconcile with some rival theories of epistemic justification. For example, on coherentism, what matters for justification is how well your beliefs cohere with each other. But coherence is not a directional notion: the beliefs that Father smokes, that Jane smokes, and that Mother smokes are equally coherent whether one infers to the first and third from the second or from the first to the second and from there to the third. It is thus unclear how coherentists can account for the difference between these two cases.

Or consider Williamson's "knowledge-first" epistemology, on which everything we know is part of our evidence, and it is possible to extend our knowledge through inductive inference (Williamson 2000, Bird 2018). If A and C are both prior to and predict B, and any other conditions for extending our knowledge through inference are met (for example, the conditional probabilities of B given A, and C given A&B, are sufficiently high), then this implies that we can start with evidence A, infer and come to know B, and on the basis of our new evidence infer and come to know C. We can, for instance, learn that Father smokes, infer and come to know that Jane smokes, and on that basis infer and come to know that Mother smokes. So Williamson's epistemology seems to imply that we can extend our knowledge through inferences the model of inference defended here deems irrational.

These implications show that questions of evidence and questions of inference are closely related. We cannot draw conclusions about what kinds of inferences are warranted without answering questions about what evidence we have, and we cannot draw conclusions about how the evidence we have supports other propositions without answering questions about inference.

Bibliography

- Audi, Robert, ed. (2015). *The Cambridge Dictionary of Philosophy, Third Edition* (Cambridge University Press).
- Bartha, Paul (2019). "Analogy and Analogical Reasoning," *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, URL = <https://plato.stanford.edu/archives/spr2019/entries/reasoning-analogy/>.
- Bhogal, Harjit (2017). "Minimal Anti-Humeanism," *Australasian Journal of Philosophy* 95: 447–460.
- Bird, Alexander (2018). "Evidence and Inference," *Philosophy and Phenomenological Research* 96: 299–317.
- Carnap, Rudolf (1950). *Logical Foundations of Probability* (University of Chicago Press).
- Climenhaga, Nevin (2017a). "How Explanation Guides Confirmation," *Philosophy of Science* 84: 359–68.
- Climenhaga, Nevin (2017b). "Inference to the Best Explanation Made Incoherent," *Journal of Philosophy* 114: 251–73.
- Climenhaga, Nevin (2020). "The Structure of Epistemic Probabilities." *Philosophical Studies* 77(11).
- Climenhaga, Nevin (ms-a). "Are Simpler Worlds More Probable?"
- Climenhaga, Nevin (ms-b). "Epistemic Probabilities are Degrees of Support, not Degrees of (Rational) Belief."
- Climenhaga, Nevin, Lane DesAutels, and Grant Ramsey (forthcoming). "Causal Inference from Noise," *Noûs*.
- Copi, Irving M., Carl Cohen, and Kenneth McMahon (2014). *Introduction to Logic: Fourteenth Edition* (Pearson Education Limited).
- Crupi, Vincenzo, Branden Fitelson, and Katya Tentori (2008). "Probability, Confirmation, and the Conjunction Fallacy," *Thinking & Reasoning* 14: 182–199.
- Douven, Igor (2017). "Abduction," *The Stanford Encyclopedia of Philosophy*, URL = <https://plato.stanford.edu/archives/sum2017/entries/abduction/>.
- Ennis, Robert H. (1968). "Enumerative Induction and Best Explanation," *Journal of Philosophy* 65: 523–529.
- Fienberg, Stephen E. (2006). "When Did Bayesian Inference Become 'Bayesian'?" *Bayesian Analysis* 1: 1–40.
- Flew, Antony, ed. (1979/1999). *A Dictionary of Philosophy: Revised Second Edition*

- (Gramercy Books).
- Gensler, Harry J. (2002). *Introduction to Logic* (Routledge).
- Goodman, Nelson (1954). *Fact, Fiction, and Forecast* (University of London: Athlone Press).
- Haack, Susan (2018). “Proof, Probability and Statistics: The Problem of Delusive Exactness,” *La Prueba en el proceso: Evidence in the Process* (Atelier, Spain): 497–521. *University of Miami Legal Studies Research Paper* no. 18-38. Available at SSRN: <https://ssrn.com/abstract=3287440>.
- Harman, Gilbert (1965). “The Inference to the Best Explanation,” *Philosophical Review* 74: 88–95.
- Harman, Gilbert (1968). “Enumerative Induction as Inference to the Best Explanation,” *Journal of Philosophy* 65: 529–533.
- Henderson, Leah (2014). “Bayesianism and Inference to the Best Explanation,” *British Journal for the Philosophy of Science* 65: 687–715.
- Henderson, Leah (2020). “The Problem of Induction,” *The Stanford Encyclopedia of Philosophy*, URL = <<https://plato.stanford.edu/archives/spr2020/entries/induction-problem/>>.
- Hesse, Mary (1966). *Models and Analogies in Science* (Notre Dame: University of Notre Dame Press).
- Howson, Colin and Peter Urbach (2006). *Scientific Reasoning: The Bayesian Approach, Third Edition* (Open Court).
- Hume, David (1748/1999). “An Enquiry Concerning Human Understanding,” ed. Tom L. Beauchamp (Oxford University Press).
- Hurley, Patrick J. (2006). *A Concise Introduction to Logic: Ninth Edition* (Wadsworth Publishing).
- Levi, Isaac (1977). “Direct Inference,” *Journal of Philosophy* 74: 5–29.
- Lipton, Peter (1993). “Is the Best Good Enough?” *Proceedings of the Aristotelian Society* 93: 89–104.
- Lipton, Peter (2004). *Inference to the Best Explanation, Second Edition* (Routledge).
- Mill, John Stuart (1843/1930), *A System of Logic* (London: Longmans-Green).
- Norton, John (2010). “A Survey of Inductive Generalization.” Retrieved from <https://philpapers.org/rec/NORASO>.
- Pearl, Judea (2000). *Causality: Models, Reasoning, and Inference* (Cambridge University

- Press).
- Pollock, John (1986). *Contemporary Theories of Knowledge* (Towota, NJ: Rowman And Littlefield Publishers).
- Reid, Thomas (1785/1895). *Essays on the Intellectual Powers of Man. The Works of Thomas Reid, Volume 3, 8th Edition*, ed. Sir William Hamilton (Edinburgh: James Thin).
- Russell, Bertrand (1948). *Human Knowledge: Its Scope and Limits* (New York: Simon and Schuster.)
- Sanford, David H. (2015). "Inference," in *The Cambridge Dictionary of Philosophy*, p. 509.
- Schaffer, Jonathan (2016). "Grounding in the Image of Causation," *Philosophical Studies* 173: 49-100.
- Schechter, Joshua (2013). "Rational Self-Doubt and the Failure of Closure," *Philosophical Studies* 163: 428–452.
- Schupbach, Jonah (2016). "Competing Explanations and Explaining-Away Arguments." *Theology and Science* 14: 256–267.
- Spirtes, Peter, Clark Glymour, and Richard Scheines (2000). *Causation, Prediction and Search, Second Edition* (M.I.T. Press).
- Swinburne, Richard (2001). *Epistemic Justification* (Oxford University Press).
- Swinburne, Richard (2004). *The Existence of God, Second Edition* (Oxford: Oxford University Press).
- Thagard, Paul (1978). "The Best Explanation: Criteria for Theory Choice," *Journal of Philosophy* 75: 76-92.
- Van Fraassen, Bas (1989). *Laws and Symmetry* (Oxford: Oxford University Press)
- Vaughn, Lewis (2009). *The Power of Critical Thinking: Third Edition* (Oxford University Press).
- Weintraub, Ruth (2013). "Induction and Inference to the Best Explanation," *Philosophical Studies* 166: 203–16.
- Williamson, Timothy (2000). *Knowledge and its Limits* (Oxford: Oxford University Press).