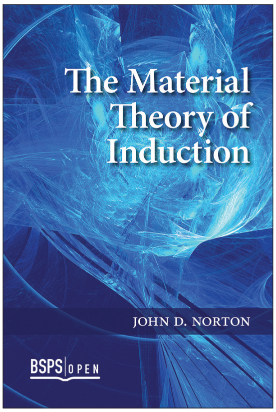

THE MATERIAL THEORY OF INDUCTION

by John D. Norton

ISBN 978-1-77385-254-6

Acknowledgement: *We acknowledge the wording around open access used by Australian publisher,* **re.press**, *and thank them for giving us permission to adapt their wording to our policy* http://www.re-press.org

1

The Material Theory of Induction Stated and Illustrated

1.1. The Terms "Induction" and "Inductive Inference"

This is a book about induction and inductive inference. Since these terms may mean different things to different people, it is worth fixing what they mean at the outset. Traditionally, induction has had a narrow meaning. At its narrowest, it refers to "induction by simple enumeration," the inference from "Some *A*s are *B*" to "All *A*s are *B*." This is an example of "ampliative inference," for we have amplified the instances to which our knowledge applies. The premise applies just to the few cases of *A*s at hand; the conclusion applies to all. I take this idea of ampliation in its most general sense to be what induction is about. I shall use "induction" and "inductive inference" as the general terms for any sort of ampliative inference. That is, they are licit inferences that lead to conclusions stronger deductively than the premises or even just conclusions that differ from those that can be inferred deductively from the premises. Therefore, the terms embrace what is sometimes called "abductive inference," which is an inference to something that explains an otherwise puzzling phenomenon.

A still broader form of induction commonly goes under the name of "confirmation theory." It typically has no inferences with premises and conclusions. Rather, it looks at degrees of support between propositions. The best-known and dominant form is probabilistic support. The conditional probability $P(H \mid E)$ represents the total inductive support an hypothesis accrues from all evidence, including our background knowledge, written as E. One then tracks how the support between hypothesis and evidence changes as the evidence changes. This form of analysis will be included under the terms "induction" and "inductive inference."

My use of the terms "inference" and "infer" will follow what I take to be the traditional usage and the one that is still most common. That is, an inference from proposition *A* to proposition *B* is a logical relation between the two propositions as sanctioned by some logic. When we infer from *A* to *B*, we merely trace through that logical relation. The usage is analogous to that of "add." When we add seven to five to arrive at twelve, we are simply tracing through the relation 5 + 7 = 12 among the three numbers as authorized by ordinary arithmetic.

This usage is to be contrasted with a psychologized notion of the term "inference" that will *not* be employed here. According to this latter usage, to say that we infer from proposition *A* to proposition *B* only records a fact of our psychology: that we proceed from a belief in *A* to a belief in *B*, without a requirement that this transition is authorized by some logic. While I understand the distinction is important to those who work in the psychology of belief, it seems to me a troublesome redefinition of a term whose normal usage is already well established. Could not another word have been found? Perhaps the redefinition is supported by the usage of the term that presupposes an agent that infers. A similar redefinition might insist that saying "we add seven to five to arrive at twelve" merely reports our belief in the summation with no supposition that it conforms with arithmetic. I would find that redefinition equally troublesome.[1]

Throughout this volume, unless some context demands an exception, I will restrict notions of inference and logic to relations of deductive and inductive support between propositions, independently of our beliefs and thought processes.

1 Harman (2002, p. 173) gives a clear statement of the psychologized notion of inference that is *not* employed in this book: "Inference and reasoning are psychological processes leading to possible changes in belief (theoretical reasoning) or possible changes in plans and intentions (practical reasoning). Implication is more directly a relation among propositions." This usage is incompatible with the longstanding and pervasive usage of "rules of inference" as designating licit manipulations and argument schemas, such as modus ponens and various syllogisms. See, for example, Boole (1854, chap. 15) and Copi (1967, p. 36 and inside back cover).

1.2. The Formal Approach to Induction

My contention is that the broad literature on induction is built on faulty foundations. It has long sought as its most basic goal to develop inductive inference as a formal system akin to deductive logic and even ordinary arithmetic. What is distinctive about these systems is that they are non-contextual, universal, and governed by simple rules. If we have six cartons of a dozen eggs each, arithmetic tells us that we have seventy-two eggs overall. It also tells us that if we have six troupes of a dozen acrobats, then we have seventy-two acrobats overall. Arithmetic tells us that when it comes to counting problems like this we can ignore almost everything except the numbers appearing in the descriptions. We extract those numbers and then see if our arithmetic provides a schema that covers them. In this case, we find in our multiplication tables that

$$6 \times 12 = 72.$$

This is really a schema that says (among other things) if you have six *groupings* of twelve *individuals*, then you have seventy-two *individuals* overall. It is a schema or template since it has empty slots, indicated by the words "grouping" and "individuals" in italics; and we generate truths about specific systems by inserting appropriate, specific terms into the slots. Insert "carton" and "egg," and we generate a numerical fact about eggs. Insert "troupe" and "acrobat," and we have a numerical fact about acrobats.

This example illustrates the key features typically sought in an inductive logic. It is to be non-contextual, universal, and formal. The numerical facts of arithmetic are non-contextual—that is, independent of the context. In abstracted form, they hold for eggs, acrobats, and every other sort of individual. The rules are universal; they do not come with restrictions to particular domains. It is the same arithmetic for eggs as for acrobats. And the rules are formal in the sense that they attend only to the form of the sentence asserting the data: six … of twelve …. The matter—eggs or acrobats—is ignored.

Deductive logic has developed similarly as a universal, non-contextual formal theory; and it enjoys extraordinary success. It has been a reasonable and attractive project to try to find a similar account of inductive inference. A universal formal theory of induction would enable us to focus attention just on the specifically inductive-logical parts, ignoring all the material complications of the much larger inductive enterprise. And we would hope eventually to generate great theorems of tremendous power and scope, perhaps rivaling those of arithmetic and deductive metalogic.

1.3. Problems of the Formal Approach

However, the formal approach is a failed project. The simple formal rules that worked so well for deductive inference have no counterpart in inductive inference. In antiquity, we were quite confident of the deductive schema

> All *A*s are *B*.
> Therefore, some *A*s are *B*.

Yet its inductive counterpart, enumerative induction—

> Some *A*s are *B*.
> Therefore, all *A*s are *B*.

—was already the subject of doubt and even ridicule in antiquity. Inductive logic never really caught up. While deductive inference has settled into the grey maturity of arcane theorem proving, inductive inference has remained an erratic child. For philosophers, the words "induction" and "problem" are routinely coupled.

There are, as we shall see later, a plethora of modern accounts of induction. But none succeed with the simple clarity of deductive logic. We should infer inductively, we are told, to the best explanation. But we are given no comparably precise account of what makes one explanation better than another—or even precisely what it is to explain something. Efforts to make these notions precise raise more problems than they solve. Elsewhere, we are told that all of inductive logic is subsumed by probability theory. Chapters 10 to 16 are devoted to arguing that the resulting theory has failed to provide a universal account of inductive inference. The probabilistic enterprise has become so many-headed that no single formula captures the difficulty. The account is sometimes too strong and imposes properties on inductive inference it should not have. It is sometimes used too permissively so that any inductive manipulation one might conceive of is somehow embraced by it. It is almost always too precise, fitting exact numbers to relations that are not that exact.

So how are we to think about inductive inference? Formal theories of induction distinguish the good inductive inferences from the bad by means of universal schemas. In their place, I propose a material theory of induction.[2] According to this view, what separates the good from the bad inductive inferences are background facts—the *matter* of the inference, as opposed to its *form*. To put it another way, we locate what authorizes an inductive inference not in some universal, formal schema but in facts that prevail in the domain of the inference.

1.4. Inductions on Crystal Forms[3]

An example will make the problems of the formal approaches clearer and the idea of a material theory of induction more concrete. We shall consider an elementary inductive inference in science that is so routine that we may even fail to notice that it is an induction. Consider a chemist who prepares a new salt of some metal and notes its particular crystalline form. It is routine for the chemist to report the form not only as the form of the particular sample but as the form of the salt generally. For crystals have quite regular properties, and crystals of different substances have characteristic differences. Nonetheless, it is an inductive inference from the one sample to all. Even if the inductive character of the inference is easily overlooked, we should expect a good treatment of it from an account of inductive inference.

To develop the example, we need to appreciate that adequate reporting of the crystalline structure of a new salt is somewhat delicate. For the individual crystals of one salt may have many different shapes. In the early history of work on crystals, it proved to be quite complicated to find a simple and robust system of classification. This complication will become a central concern of the material analysis of these inductive inferences.

---

2 For earlier accounts, see Norton (2003, 2005).

3 My thanks to Pat Corvini for correcting errors in an earlier version of this section and also Section 1.9 below; and later for providing an extensive list of typographical errors in the Prolog and Chapters 1 and 2.

Crystallographic analysis now categorizes crystal forms according to the axes characteristic of the shape. The simplest of the seven crystallographic systems is the *cubic* or *regular system*. The crystals of common table salt, sodium chloride, fall into this system. It is characterized by three perpendicular axes of equal length. A cube conforms to this system; it takes no great geometrical insight to see that a cube has these three perpendicular axes of equal length. The same is true of a regular octahedron, which also conforms to the system. Sodium chloride normally crystallizes in cubes. However, in special environments, such as in the presence of urea, it can crystalize as octahedra.

One might imagine that the cube and octahedron are the only shapes that crystals in the cubic system can adopt. Matters are more complicated, however, for there are many more shapes in this system. The mineral spinel lies within the cubic family and forms octahedral crystals. However, spinel can also form many misshapen octahedral crystals, as shown in Figure 1.1.

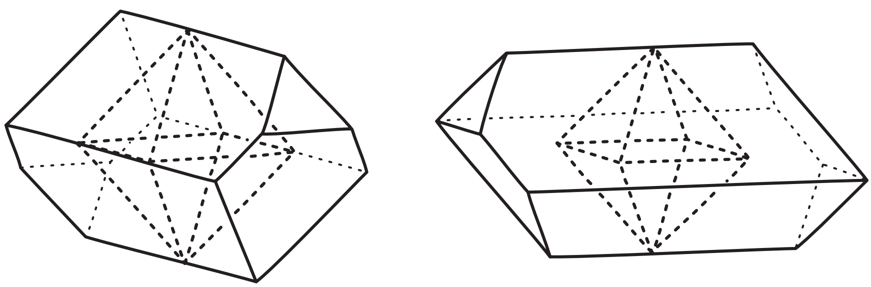

Figure 1.1. Misshapen octahedra.[4]

The octahedral character of the crystals arises from their faces being parallel to those of a fictional regular octahedron, which we might imagine secretly buried within the crystal.

Crystals have natural cleavage planes. A crystal cube of sodium chloride will cleave along planes parallel to the cube's surfaces. The mineral fluorspar represents an unusual case. Although it is in the cubic family

4 Illustration based on Miers (1902, p. 11, Figs. 9 and 10).

and crystalizes in cubes, it cleaves along planes that eventually expose an octahedral shape. Figure 1.2 shows successive cleavages.

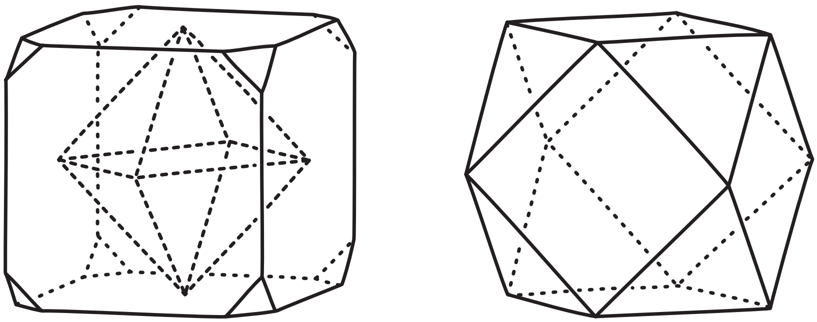

Figure 1.2. Cleaving fluorspar.[5]

In the process of cleavage, we pass through many more complicated cube shapes with corners removed to different extents. The shape on the right of Figure 1.2 is such an intermediate form. These multi-faceted shapes and many more are licit forms for certain crystalline substances within the cubic system.

All of these shapes are different from the crystalline shapes permitted to barium chloride, for barium chloride is monoclinic. This means that its crystals are characterized by three unequal axes, two of which intersect at an oblique angle, and a third that is perpendicular to them. Instead of a cube, its primitive form—the simplest crystal shape—is a right prism with a parallelogram base. This is shown in Figure 1.3, where the parallelogram is the rearmost face. Alternatively, one may generate the shape by starting with a right prism with a rectangular base and inclining it to one side (hence “mono-cline”). In Figure 1.3, the inclination is towards the right of the figure.

5 Illustration based on Miers (1902, p. 14, Figs. 17 and 18).

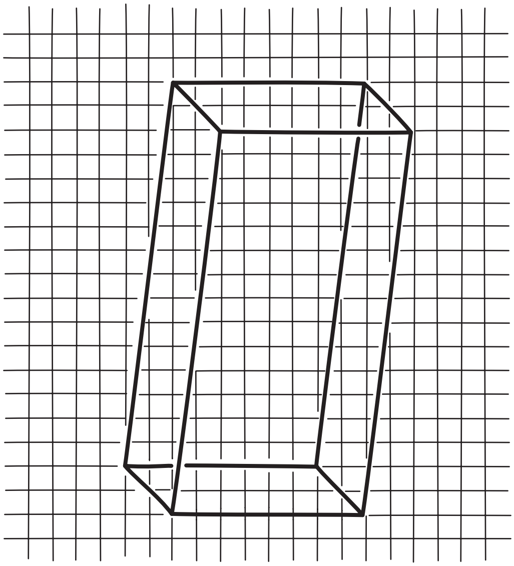

Figure 1.3. Primitive form of the monoclinic system.

The range of crystal shapes allowed in the monoclinic system is related to this form in the same way that those allowed in the cubic system are related to a cube.

When a new metallic salt is prepared, the chemist will simply assert that such-and-such is the form of the salt's crystals. This is an inductive inference and one of breathtaking scope. On the strength of just a few samples, the chemist is quite prepared to infer the crystal system of all samples of the salt:

> This sample of salt *A* belongs to crystallographic system *B*.
> Therefore, all samples of salt *A* belong to crystallographic system *B*.

1.5. Curie and Radium

Perhaps the most famous of all episodes in crystal formation was Marie Curie's separation of radium from uranium ore by fractional crystallization. The massive labor of extracting radium from the pitchblende ore is the stuff of scientific legends, Nobel Prizes, and a 1943 MGM movie. The radioactive elements—polonium, radium, and actinium—exist in such trace quantities that several tons of uranium ore residue had to be treated to recover just a few decigrams of radium. A decigram, a tenth of a gram, is a mere speck. The process of recovering the radium was arduous. From each ton of ore, after much processing, about eight kilograms of barium chloride was recovered. Radium chloride is present in barium chloride as a trace impurity. Radium's presence is revealed by its great radioactivity.

The final separation of the radium chloride from the barium chloride was difficult to achieve since radium and barium behave in similar ways chemically. The separation depends on the fact that radium chloride is less soluble in water than barium chloride. If the barium chloride in solution is concentrated by boiling and cooling until it forms crystals, the crystals will harbor more radium chloride. The solution remaining above the crystals had one fifth of the radioactivity of the original, Curie reported. While that seems like a large increase, the quantity of radium present in the crystals was so tiny that it fell far short of what was required for substantial separation. Curie needed to repeat the process over and over: redissolving and recrystallizing to form more fractions, recombining them according to their radioactivity, and doing it again and again. In all, she needed to carry out several thousand crystallizations.

All of this is described in her doctoral dissertation (Curie 1904), presented to the Faculté des Sciences de Paris in June 1903. There, she reported on the analytic work carried out in the few years before with her husband, Pierre Curie. The feature of the radium chloride that attracted most attention was its powerful radioactivity. In spite of the thousands of crystallizations performed, the crystallographic properties of radium chloride barely rated a mention. In the ninety-four pages of the dissertation, there are only a few complete sentences on the crystallographic form, and they bleed off into less certain reports on the colors of the crystals that, she suspected, would prove of practical use in the separation:

> The crystals, which form in very acid solution, are elongated needles, those of barium chloride having exactly the same appearance as those of radium chloride. Both show double refraction. Crystals of barium chloride containing radium are colourless, but when the proportion of radium

> becomes greater, they have a yellow colouration after some hours, verging on orange, and sometimes a beautiful pink. This colour disappears in solution. Crystals of pure radium chloride are not coloured, so that the colouration appears to be due to the mixture of radium and barium. The maximum colouration is obtained for a certain degree of radium present, and this fact serves to check the progress of the fractionation.
>
> I have sometimes noticed that formation of a deposit composed of crystals of which one part remained uncoloured, whilst the other was coloured, and it seems possible that the colourless crystals might be sorted out. (Curie 1904, p. 26)

Curie and soon others separated out only minuscule quantities of radium. Yet that radium chloride forms crystals just like those of barium chloride entered the literature quite quickly. In his 1913 survey of radioactive substances, Ernest Rutherford reported:

> Radium salts crystallise in exactly the same form as the corresponding salts of barium. The crystals of radiferous barium chloride several hours after preparation usually assume a yellow or rose tint. The intensity of this colouration depends on the relative proportions of barium and radium present in the crystal. Nearly pure radium chloride crystals do not show this colouration, indicating that the presence of barium is necessary. (Rutherford 1913, p. 470)

The facts are reported as having quite general scope, even though the instances of observed radium chloride crystals must have been few, given the enormous labor needed to create them in tiny quantities. Nonetheless, both Curie and Rutherford seemed quite certain of the generalization. Rutherford's report looks like little more than a paraphrase of Curie's remark.

1.6. A Formal Analysis

If we approach inductive inference formally, how are we to accommodate this induction? We need only investigate a few simple formal attempts to see just how poor the formal analysis is. The inference looks like a type of enumerative induction with the schema

> Some (few) *A*s are *B*.
> Therefore, all *A*s are *B*.

Yet this alone cannot be what authorizes the induction. For almost every substitution of the *A*s and *B*s would yield a feeble induction. To get an induction of the strength seen by Curie and Rutherford, we have to be selective in what is substituted for *A* and *B*. The *A*s have to be specific chemical types, such as radium chloride or barium chloride, as opposed to the hundred and one other types of stuff that Curie found in her vats. More importantly, the induction works only for carefully chosen properties of *B*. There are many ways of describing crystal forms. Virtually none of them support a strong inductive inference.

To revert to the simpler example, one may find that some particular crystal of common salt is a perfect cube. However, no chemist would risk the induction to all crystals of common salt having exactly that shape. It was only after serviceable systems of crystallography were introduced that the right property was found. Individual crystals of common salt fall into the cubic or regular system, and this property can be inserted into the schema of enumerative induction to form the generalization.

The problem of finding the right descriptions challenged generations of crystallographers. Indeed, for a long time, many held that crystal forms admit no simple systematization so that exactly this sort of induction would be denied. The scientist, historian, and philosopher of science William Whewell gave a lively account of these hesitations—and of how Romé de l'Isle and René Just Haüy after 1780 sought to resolve the problems—in his *History of the Inductive Sciences* (1837, vol. 3, book 15, chaps. 1–2).

These difficulties make it a matter of some delicacy to specify in formal terms just what property of the radium chloride crystals can be generalized. Curie and Rutherford used parasitic locutions: the crystals of radium chloride are the same as those of barium chloride. Hence, Marie Curie in her 1911 Nobel Prize address chose a technical locution to describe the crystal form of radium chloride: "In chemical terms radium differs little from barium; the salts of these two elements are isomorphic, while those of radium are usually less soluble than the barium salts" (Curie [1911] 1999). Isomorphism is a term of art used then and now to describe the circumstance in which two different substances have very close chemical and crystalline properties (see Miers 1902, p. 213). Curie's use of the term saved her the need of describing in more detail the precise structure possessed by the salts of radium. It was familiar knowledge for chemists that barium chloride has such-and-such a monoclinic crystalline form. The declaration of isomorphism indicated that radium chloride had this form too.

If the schema of enumerative induction is to function as a general logic, the restrictions on just what may be substituted for *A* and *B* have to be abstracted, regularized, and formalized, and then included in the schema. The problem is that the restrictions that must be added are so specific that one despairs of finding a general formulation. Presumably, a general logic cannot append clauses of the form: "If *A* is a substance that manifests in crystalline form, then *B* must be one of the known crystal forms as sanctioned by modern crystallography." This is a little short of offering a huge list in which we inventory the specific inferences that are allowed. This would not be a logic but a catalog whose guiding rationale would be hidden.

A more promising approach is to draw on a popular philosophical notion devised for this sort of application: we require that *A* and *B* must be natural-kind terms. These are terms adapted to the divisions arising in nature ("is crystallographically regular"), as opposed to artificial divisions introduced by humans ("looks like a cubist sculpture"). The hope is that we succeed in delimiting the good inductive inferences by restricting the schema explicitly to natural-kind terms.

The approach fails at multiple levels. First, it fails because the good inductions on crystal forms are still narrower. It is surely a natural-kind term for a crystal to be a perfect cube, one of the five Platonic solids. Yet an induction on common salt that uses the property fails to be a good induction by the standards of the crystallographers. Second, the schema is only viable if one can give a general formula that specifies what a natural-kind term is. A common characterization of natural-kind terms is that they support induction (Bird and Tobin 2010, sec. 1.1). This means that we are allowed to generalize relations found in a few cases to hold between natural-kind terms. If we append this characterization of natural-kind terms to the schema of enumerative induction, the schema is rendered circular. For to require that the schema can only be used on terms *A* and *B* that support induction is just a fancy way of saying that the schema only works when it works. Another common characterization of natural-kind terms is that they appear in natural laws. If we try to include this characterization in the specification of the schema, we face similar circularities when we try to state just what we mean by "law." Are they true relations that obtain between natural kinds?

1.7. A Bayesian Attempt[6]

The preceding section sought to develop the simple schema of enumerative induction to convert it into a serviceable schema with universal application. The efforts were unsuccessful. Might a different approach that employs probabilistic analysis fare better? What if we seek help from Bayesian analysis? We seek a vindication of the inference from "Some (few) *A*s are *B*" to "All *A*s are *B*" that relies essentially on the probabilistic character of relations of support. It should not merely adopt antecedently some version of the idea that the proposition "All *A*s are *B*" accrues support from the proposition that "Some *A*s are *B*" and then just restate it in probabilistic language. We saw that it was precisely this idea that proved unsustainable in the last section. Simply translating the idea into probabilistic language would only serve to hide the difficulties behind a veil of numbers and formulae. In addition, we should like the probabilistic analysis to show us that "Some (few) *A*s are *B*" can provide strong support for "All *A*s are *B*."

There are many ways that one can give Bayesian analyses of this problem. Let me sketch just one. We write *H* for the hypothesis that a newly prepared salt belongs to some particular crystallographic system. We

6 I thank Nick Huggett for helping me to think through revisions to this section and the next.

write E for the evidence that a number of samples are each observed to belong to that class. If there are n samples, we can write $E = E_1 \& E_2 \& \ldots \& E_n$, where E_i asserts the evidence in the ith case. The probability of interest is $P(H \mid E)$, the probability of the hypothesis H given the evidence E. This represents the inductive support afforded to H by E if we think of the probabilities objectively. Or, if we interpret the probabilities subjectively, it is the belief that we have in H given that we know E. We are interested in seeing how the posterior probability $P(H \mid E)$ compares with the prior probability $P(H)$; that is, we seek to determine how the probability of H changes when we incorporate our learning of evidence E. These changes will tell us the evidential import of E. An increase in probability is favorable evidence; a decrease is unfavorable.

We can compute these changes by means of Bayes' celebrated theorem. In a form suitable for this application, it asserts

$$\frac{P(H \mid E)}{P(\sim H \mid E)} = \frac{P(E \mid H)}{P(E \mid \sim H)} \frac{P(H)}{P(\sim H)}.$$

We will not compute $P(H \mid E)$ directly but only how incorporating E alters the balance of probability between the hypothesis H and its negation $\sim H$. That is, we can see how the ratio of the prior probabilities $P(H)/P(\sim H)$ changes to $P(H \mid E)/P(\sim H \mid E) = r$. From this last ratio, $P(H \mid E)$ can be recovered as

$$P(H \mid E) = \frac{r}{r+1}.$$

Bayes' theorem tells us that the controlling quantities are the two likelihoods $P(E \mid H)$ and $P(E \mid \sim H)$. The first is easy to compute. It expresses the probability that we have the evidence E if the hypothesis H is true. The hypothesis H says that all samples belong to a particular crystallographic system. Hence, the n samples at hand must belong to that system. So the probability is unity that we have evidence E: $P(E \mid H) = 1$.

The other likelihood, $P(E \mid \sim H)$, is much harder to determine. It requires us to assess the probability of the evidence if the hypothesis is false. Determining this quantity requires some creative imagination, for we have no precise prescription for how the hypothesis might fail. The likelihood will vary depending on how we judge the hypothesis might fail. If the only possibility for failure is that the salt belongs to one of the other crystallographic classes, then there is no possibility of the evidence E obtaining. Then $P(E \mid {\sim}H) = 0$. Inserting this into Bayes' theorem leads to $P(E \mid H) = 1$; the hypothesis is maximally probable.

But things are more complicated. E can be reported if there are observational errors so that the evidence is misreported. Or it may turn out that the salt is dimorphous or even polymorphous. This means that the salt can crystallize into two or more systems. So there is some chance—perhaps small, perhaps large—that the evidence E_i obtains, even if H is false.

We will set these concerns aside. Let us set the probability to q so that $P(E_i \mid {\sim}H) = q$ and suppose that each of the samples is taken under independent conditions with the supposition of the falsity of H. Then, obtaining each E_i is probabilistically independent of the others, and the probability of the conjunction is just a simple product of terms:

$$P(E \mid {\sim}H) = P(E_1 \,\&\, E_2 \,\&\, \ldots \,\&\, E_n \mid {\sim}H) = P(E_1 \mid {\sim}H) \cdot P(E_2 \mid {\sim}H) \cdot \ldots \cdot P(E_n \mid {\sim}H) = q^n$$

Bayes' theorem now becomes

$$\frac{P(H \mid E)}{P({\sim}H \mid E)} = \frac{1}{q^n}\frac{P(H)}{P({\sim}H)}.$$

Here we have a nice limit result. As n becomes large, q^n can be brought arbitrarily close to 0, as long as $q < 1$. Hence, the ratio of likelihoods $1/q^n$ becomes arbitrarily large, so that the ratio $r = P(H \mid E)/P({\sim}H \mid E)$ also grows arbitrarily large. This corresponds to the posterior $P(H \mid E) = r/(r + 1)$ coming arbitrarily close to unity. And this means that the support for or belief in H approaches certainty. This limiting result is comforting, for it means that we do not need to worry about the particular values that we might assign to the priors. Whatever influence their values may have had on the final result is "washed out" by the limit process. This is for the best since the prior probabilities $P(H)$ and $P({\sim}H)$ would have to be plucked from the air.

1.8. Where the Bayesian Analysis Fails

If one inclines to numerical and algebraic thinking, the foregoing may seem like a very satisfactory analysis. It has brought mathematical precision to what at first seemed like an intractable problem. There is even a little limit theorem in which priors are washed out. All that is an illusion. There are few if any gains in the analysis. And the harm done is great since we have convinced ourselves that we have solved a great problem when we have not. Any positive result achieved has little to do with the probabilistic properties supposed for relations of inductive support and everything to do with the choices we make externally to the analysis. We shall see that the long-term results are determined by our antecedent choice of prior probabilities, which prove to be narrowly constrained to two extreme, dogmatic possibilities. The short-term results depend critically on arbitrarily chosen numbers. Finally, the necessary condition for any successful result lies in choosing a description of the hypotheses and evidence that is delicately tuned to the properties of the system. Without such a description, inductive success is impossible. With it, success is assured for virtually any approach.

1.8.1. External Inductive Content

The first problem is that the analysis is heavily dependent on judgments of probability that are supplied externally to the analysis. That is, we must set prior probabilities that presume either a dogmatic skepticism or an unreasonable credulity concerning the universal hypothesis *H*. There is no other option.

To avoid the danger of these externally specified assumptions prejudging the result, we might require a prior probabilistic independence of the individual items of evidence, E_1, E_2, …, E_n. This avoids an antecedent assumption of them being connected by the universal hypothesis *H*. That is, we would have

$$P(E) = P(E_1 \;\&\; E_2 \;\&\; \ldots \;\&\; E_n) = P(E_1)\, P(E_2) \;\&\; \ldots \;\&\; P(E_n) = s^n,$$

where for simplicity I have assumed an equal probability $0 < s < 1$ for each $P(E_i)$. The result is immediately disastrous. A version of Bayes' theorem now tells us that

$$P(H \mid E) = \frac{P(E \mid H)}{P(E)} P(H) = \frac{1}{s^n} P(H).$$

As the number of instances of n increases, s^n decreases and can be brought arbitrarily close to zero, which means that $1/s^n$ can be made arbitrarily large. Since $P(E \mid H)$ can never exceed unity, probabilistic consistency requires that we can no longer choose our prior probability $P(H)$ freely. We must have $P(H) \leq s^n$. Since s^n can be brought arbitrarily close to zero when n is large enough, we must somehow choose a prior probability $P(H)$ close enough to zero that anticipates in advance the number of items of evidence that may appear. The only secure value is a zero prior probability: $P(H) = 0$. In this worst case, we preclude learning from the evidence, since $P(H) = 0$ forces $P(H \mid E) = 0$ no matter what evidence E is presented. We must commit to a prior skepticism about the universal hypothesis H.

It is entirely reasonable to respond that this shows that presuming prior probabilistic independence of the individual items of evidence E_1, E_2, …, E_n is not benign after all. The assumption of independence encodes a dogmatic skepticism concerning the universal hypothesis H. But the alternative is equally troublesome. If we now admit the possibility of a prior probabilistic dependence among the items of evidence, we commit to unreasonable credulity concerning the universal hypothesis H. Here is why:

To avoid prior skepticism about H, we must free ourselves of the need to set $P(H)$ arbitrarily close to zero. We do this by ensuring that $P(E) = P(E_1 \,\&\, E_2 \,\&\, \ldots \,\&\, E_n)$ does not become arbitrarily small as n grows large. We expand $P(E)$ as

$$\begin{aligned} P(E) &= P(E_1 \,\&\, E_2 \,\&\, \ldots \,\&\, E_n) \\ &= P(E_n \mid E_1 \,\&\, E_2 \,\&\, \ldots \,\&\, E_{n-1})\, P(E_{n-1} \mid E_1 \,\&\, E_2 \,\&\, \ldots \,\&\, E_{n-2}) \ldots P(E_2 \mid E_1)\, P(E_1). \end{aligned}$$

We preclude $P(E)$ becoming arbitrarily small by requiring that $P(E_n \mid E_1 \,\&\, E_2 \,\&\, \ldots \,\&\, E_{n-1})$ approaches unity in the limit as n grows large. This requirement says that conditioning on the evidence E_1, E_2, …, E_{n-1} requires the limiting probability of E_n to be arbitrarily close to unity. This is close to assuming H itself. For informally it says that being an instance of H is projectable in this sense: if we have seen $n - 1$ instances of H with increasing n, we approach probabilistic certainty that the next, nth item will also be an instance with H.

The credulity toward H lies in the permissiveness of this result. It turns out that we approach probabilistic certainty not just for the next instance of H, but for the next N instances of H after it—no matter how large N is. For a simple variant of the last calculation shows that the conditional probability

$$P(E_n \& E_{n+1} \& \ldots \& E_{n+N} \mid E_1 \& E_2 \& \ldots \& E_{n-1})$$

must also approach unity as n and N grow large. Our confidence in projectability is not limited just to the universal hypothesis H, but to any hypothesis of which the items of evidence are an instance, no matter how curious the hypotheses. The hypothesis may be that all samples of radium chloride were prepared by Curie; or that all are in Paris; or that all are in the northern hemisphere.

In sum, we cannot simply present the evidence as bare data and have the Bayesian analysis tell us its import. We have to add prior probabilities and there is no benign way to set them. We must choose antecedently between those that commit us to a dogmatic skepticism or to an unreasonable credulity. This difficulty of Bayesian analysis has long been recognized.[7] Richard Jeffrey (1983, p. 194) was sufficiently disturbed by it that he concluded "willingness to attribute positive [prior] probability to a universal generalization is tantamount to willingness to learn from experience at so great a rate as to tempt one to speak of 'jumping to conclusions.'" This example illustrates a quite general result reviewed in Chapter 12: formal analyses within a calculus of inductive inference cannot be freed from their dependence on externally supplied inductive content.

7 For a brief review, see Norton (2011, pp. 430–31).

1.8.2. Curie Did Not Take a Large *n* Limit

The second issue is that the analysis has solved the wrong problem. Curie was sure of the result already from just a few samples. She did not need to look at *n* samples and ponder the result as this *n* grew arbitrarily large. This "small *n*" result can be addressed in the Bayesian system, but it requires us to insert numbers. We need concrete values for q and for the priors $P(H)$ and $P(\sim H)$ in order to determine whether the analysis supports Curie's analysis. Which are the right values? Can we find them? Or are our selections just hunches driven by dim feelings of what is reasonable?

We now must face the awkward problem of all Bayesian analysis: namely, that it introduces specific probability numbers while no such numbers are in evidence in the inductive practice. Just which value is appropriate for $P(E_i \mid \sim H)$? Is it 0.1? Or 0.5? What of the prior probabilities? If we think of the probabilities as measuring objective degrees of support, then we have no good basis for assigning the prior probabilities, and the whole small *n* calculation will rest on a fabrication. If we think of probabilities subjectively so that they merely reflect our freely chosen opinion, we are no better off. The hope, in this case, is that the accumulation of evidence will wash out the individual prejudices we introduced by arbitrary stipulation of our prior belief. This washing out does not happen precisely because we are limited to the small *n* analysis.

More generally, this "solving the wrong problem" is an infraction committed repeatedly in Bayesian analyses. There are a few simple computations that serve as examplars, and the exercise in Bayesian analysis is to modify the problem actually posed in successive steps until it resembles one of them. In this case, the original problem is transformed into the problem of distinguishing a double-headed coin (hypothesis H) from a coin that has probability q of showing a heads (hypothesis $\sim H$). We are given the evidence E of n independent tosses, all of which show heads.

These first two problems are familiar and generally addressed by making the analysis more complicated. If selecting appropriate likelihoods or prior probabilities is troublesome, then a skeptical reader may be reassured that further Bayesian analysis will surely vindicate exactly the selections needed to get the result promised. My prediction, however, is that this maneuver will not solve the problem. It will merely enlarge the analysis and exile such problems to remote corners, where they will proliferate. The problems will just be harder to see because the analysis will have become so much more complicated.

1.8.3. Finding the Right Description

The third problem is, in my view, the most serious. The Bayesian analysis began by declaring the hypothesis that the salt has crystals belonging to a certain crystallographic system and that the observed instances all conformed to this system. Once this description is given, the most important part of the inductive analysis is over; for once we know that these are the terms in which the problem should be described, then almost any analysis will succeed. Enumerative induction will quickly return something like Curie's result. Or, looking ahead to other accounts of induction, we can declare the evidence a severe test of the hypothesis; or best explained by the hypothesis.

Until we are able to describe things in these terms, no analysis will work, not even the Bayesian. The alternative descriptions will either be too coarse or too fine. If they are too coarse, the sorts of hypotheses investigated and affirmed under Bayesian analysis will likely end up as banal. We may affirm that radium chloride forms crystals, for example. If the descriptions are too fine, we will likely find that no hypothesis is well supported by the evidence. If, for example, we give too detailed a description of the crystal form, then the several cases at hand will differ sufficiently such that no single description fits and we will be left without a compatible hypothesis to set for H in the analysis.

The damage done by the Bayesian analysis is that it obscures exactly the most important part of the inductive analysis with a smokescreen of numbers and theorems. The essential part of the analysis is the recognition that the hypothesis and the evidence need to be described in terms of a narrow and hard-won vocabulary of crystallographic theory. The elaborate computations of Bayesian analysis mislead us into thinking that inductive problems are solved by manipulating probabilities and by proving theorems in the probability calculus. It is a seductive aura of precision that is to be resisted if we are to understand inductive inference.

It is widely acknowledged that the real challenge lies in finding the appropriate system of classification. In introducing crystallography as a “classificatory science,” William Whewell stressed that finding this appropriate description is the object of the science:

> Our classification of objects must be made consistent and systematic, in order to be scientific; we must discover marks and characters, properties and conditions, which are constant in their occurrence and relations; we must form our classes, we must impose our names, according to such marks. We can thus, and thus alone, arrive at that precise, certain, and systematic knowledge, which we seek; that is, at science. The object, then, of the classificatory sciences is to obtain *fixed characters* of the kinds of things; the criterion of the fitness of names is, that *they make general propositions possible*. (1837, pp. 212–13; emphasis in original)

Finding the right system of classification is what makes generalization possible.[8]

1.9. A Material Analysis

Formal analysis presumes that one isolates the transition from knowledge of a single case to all cases as a problem in inductive logic, and that we establish the cogency of the transition by displaying its conformity with formal principles. For example, we might seek to show conformity of the transition with an abstract schema of enumerative induction or, in the probabilistic case, with Bayes’ theorem. Hence, the inference from a single sample to all is immediately beset with the familiar problems that have troubled induction for millennia. They sustain the weary sense among philosophers that induction, trouble, and woe all go together.

8 Looking ahead, a probabilistic analysis could avail itself of the “Weakened Haüy’s Principle” (discussed below), which I argue warrants the inference materially. The analysis would derive directly from the principle that there is a high probability that all samples of radium chloride crystals are monoclinic, conditioned on the fact that Curie’s few samples are monoclinic. This is merely a probabilistic restatement of the final result already achieved. Probabilistic analysis has added nothing beyond the illusion of quantitative precision.

Chemists at the start of the twentieth century, pondering the crystalline structure of matter, would likely not have sensed that their passage from one sample to all was problematic. Indeed, they are unlikely to have thought of it in the abstract terms of theories of inductive inference at all. The century before had seen vigorous investigation into the question of just how properly to characterize the crystalline forms so that the passage from properties of one sample to all may be effected. Curie and Rutherford, if called on to defend this transition, would not have recited passages from logic books. They would have pointed to background knowledge then shared by all competent chemists.

The foundations of the successful approach to crystallographic categorization were laid by René Just Haüy in the late eighteenth and early nineteenth century. His approach was based on the idea that each distinct substance that forms crystals is built up from many, primitive geometrical nuclei, all of the same geometric shape. The mineral galena, in this theory, is built from minute cubes. In his treatise published at the time Curie was working on radium, Henry Miers (1902, p. 21) illustrated Haüy's account as in Figure 1.4:

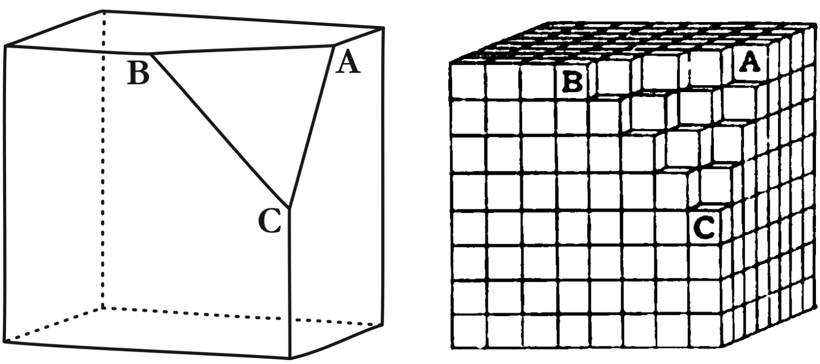

Figure 1.4. Haüy's account of crystalline shapes.[9]

9 The figure on the left is based on Miers' Fig. 38 and the figure on the right is a reproduction of Miers' Fig. 37.

The oblique face ABC of a galena crystal in Figure 1.4 is, at the smallest scale, really many staircases of these cubes. But the scale is so small that we perceive a perfectly smooth surface.

An account by a contemporary of Haüy's, Frederick Accum (1813, p. 110), summarized the theory: "He [Haüy] has also shewn that *all* crystals, however complicated their form may be, contain within them a primitive geometrical nucleus, which has an invariable form in each chemical species of crystallisable material."[10] From this theory came the essential result that every substance was characterized by a unique primitive form:

> The diversity of primitive forms ought therefore to be regarded as a certain indication of a difference in nature between two substances and the identity of primitive form indicates identity of composition, unless the nucleus is one of those solids which have a marked character of regularity; such as the cube, the regular octahedron, &c. (p. 117)

The essential qualification is that sometimes two substances may be composed of nuclei of the same form; this was likely to happen for crystals built from regular solids, like cubes. This was a quite essential qualification since Accum could list numerous cases of substances with the same crystalline form. For example, he listed ten substances based on the cube (1813, p. liv), among which were native gold, native silver, native copper, gray cobalt ore, leucite, common salt, galena, and iron pyrites.

A century later, Haüy's system had received multiple adjustments and his basic supposition was commonly bowdlerized:

> The Abbë Reny Just Hauy [*sic*], whom Dr Tutton designates the "father of modern crystallography," has enunciated the great principle that to every specific substance of definite chemical composition capable of existing in the solid condition there appears a crystallizing form peculiar to and characteristics of that substance. (Anon, p. 365)

10 This account is more succinct than Haüy's own synopsis (cf. Haüy 1807, pp. 86–101).

The view outlined was not so much a principle as a simple consequence of Haüy's theory, which, according to Accum, did not insist that each crystalline substance had its own "peculiar"—that is, unique—form.

For our purposes, the essential point is that if a chemist were to accept Haüy's theory, then one good sample of a crystalline substance would be sufficient to identify the crystallographic system to which all crystals of that substance belong. We have the following inference:

> Each crystalline substance has a single characteristic crystallographic form (Haüy's Principle).
> The sample of salt *A* has crystallographic form *B*.
> Therefore, (deductively) all samples of salt *A* have crystallographic form *B*.

This is the crudest version of how chemists pass from a single sample to all. What is notable is that it is not an inductive inference at all. The inference is deductive and authorized by early crystallographic theory.

Of course, this is an extreme case and a purely deductive inference was possible only during a brief window of a few decades during the early years of Haüy's crystallographic theory. The theory soon encountered anomalies. The shapes Haüy postulated for his nuclei could not always be stacked so as to properly fill space. Whewell (1837, p. 235) reported the collapse of Haüy's physical theory: "and when Haüy, pressed by this difficulty, as in the case of fluor-spar, put his integrant molecules together, touching by the edges only, his method became an empty geometrical diagram, with no physical meaning." A still more serious problem was the recognition mentioned above that one crystalline substance may form crystals that belong to two, three, or more crystallographic systems—called "dimorphism," "trimorphism," and "polymorphism," respectively. It was not clear how merely stacking nuclei of the same shape could yield these different shapes. Mineralogy texts of the early twentieth century routinely reported examples. William Ford's list is presented as something of a reminder of what everyone supposedly knew, rather than as a surprising novelty:

> Carbon in the forms of graphite and diamond, calcium carbonate as calcite and aragonite, iron sulphide as pyrite and

> marcasite, are familiar examples of dimorphism. The two minerals in each case differ from each other in such physical properties as crystallization, hardness, specific gravity, color, reactions with acids, etc. Titanium oxide, TiO_2, is trimorphous, since it occurs in the three distinct minerals, rutile, octahedrite and brookite. (1912, p. 80)

This means that Haüy's principle of the earlier deduction was not true, for there were cases of one substance routinely manifesting in several different crystalline forms.

But the idea of a strict regularity in the crystal forms manifested by one substance remained. So we might render a corrected version of the earlier inference accordingly:

> *Generally*, each crystalline substance has a single characteristic crystallographic form (Weakened Haüy's Principle).
> The sample of salt *A* has crystallographic form *B*.
> Therefore, (*inductively*) all samples of salt *A* have crystallographic form *B*.

We now have an inductive inference. The warranting principle is what I have called the "Weakened Haüy's Principle." What makes it inductive is the word "generally." It licenses us to proceed from one sample to all, but not with certainty.

One might imagine that this "generally" is, finally, a manifestation of some universal inductive logic. Its schema might be represented as

> Generally, *X*.
> Therefore, *X* in this case.

While we may find many instances of propositions of the form "Generally, …," they are not manifestations of a unique inductive logic. In each case, the word "generally" will have a meaning peculiar to the context. In this case, "generally" means "in so far as polymorphism does not interfere." So the nature of the risk one takes in accepting the conclusion will differ with each context.[11]

This is one illustration of how background knowledge drives inductive inferences and how such background knowledge is deeply entangled with inductive practices. Once one knows to look for it, the extent of the entanglement is quite profound. Another notion that was well established in Curie's day was isomorphism, mentioned earlier. This was then defined more precisely by Ford (1912, p. 79) as "A series of compounds which have analogous chemical compositions and closely similar crystal forms are said to make an isomorphous group." An early case of initially unrecognized isomorphism became a celebrated triumph of crystallographic analysis. Whewell (1837, pp. 226–28) reports confusion over the crystalline substance "heavy spar." Haüy found that its cleavage angles varied by three and a half degrees, depending on the origin of sample. One sample was from Sicily and one from Derbyshire. The variation was a great perplexity and a dire threat to Haüy's theory since the same nuclei could not accommodate even such a small change of angle. It turned out that the two samples were of different substances. The Sicilian sample was barium sulphate and the one from Derbyshire was strontium sulphate. Barium and strontium are both alkaline earth metals in the same column of the periodic table and have similar chemistry. They also form crystals that are very similar, although—crucially—not perfectly identical. This is a classic case of isomorphism.

When Curie remarked that the radium chloride formed crystals with "exactly the same appearance" as barium chloride, it would have been with full knowledge that the chemistry of radium mimicked closely that of barium. Indeed, that mimicry is what made the separation of the two so difficult. Hence, the familiar idea of isomorphism would have indicated that the crystals of the two chlorides should be similar. All that was really

11 While the inferences may look formally similar, they will be quite different if applied to crystals or to astronomy. Take the following proposition: "Generally, orbiting objects in our solar system orbit in the same direction as the earth." From this, we may infer with a small risk that a recently discovered asteroid will orbit in the same direction as the other objects in our solar system. The risk we take is different from that taken in crystallography. We risk the possibility that this asteroid was not formed by the same processes that formed most other objects in our solar system.

left to affirm was how close the similarity would be. It was, Curie found, “exactly the same.”

Immediately after Curie’s work was published, the chemical and crystallographic similarity of radium and barium was immediately investigated and affirmed. Runge and Precht (1903) used spectrographic and atomic weight measurements to locate radium with the other alkaline earth metals, magnesium, calcium, strontium, and barium. The expected similarity of crystalline forms was found by direct measurement of the bromides of barium and radium. As Frederick Soddy reported,

> F. Rinne … has published a careful comparison of the crystallographic relation between the bromides of radium and barium and has shown that radium bromide crystallises in the monoclinic system and is isomorphous with and crystallographically closely related to barium bromide. (1907, p. 332)

To report the isomorphism of barium and radium became standard in the literature.

We can now appreciate the great subtlety of Curie’s inference. As long as the background theories of crystallography are to be trusted, the possibility of polymorphism was the principal risk taken in generalizing the crystalline form of radium chloride from one sample to many. Hence, Curie and Rutherford were quite sanguine to report the radium salts’ crystalline form as an isomorphism with barium salts. For if there had been any polymorphism of the radium salt, they could reasonably expect a similar polymorphism to arise with the barium salt. So, with or without polymorphism, their result would stand. With that canny formulation, the result could be asserted with the confidence they showed. The only real danger was a failure of the isomorphism and, given the multiple points of agreement between barium and radium, that was easy to discount.

Let us take stock. Our starting point was a simple inductive inference from a few crystal samples to all samples. It is the sort of simple induction that should be explicated easily by an inductive logic. In particular, we would expect the logical analysis to tell us why this particular inference from “some” to “all” is so strong as to be essentially unquestioned. On closer inspection, we found appearances to be deceptive. The strength of the passage from “some” to “all” in this particular case had little to do with issues identifiable by a formal logic. It had all to do with background chemical knowledge. The confidence the chemists had for the inference resulted from the care with which Curie and Rutherford located the inference within a complicated network of chemical ideas that had been devised over the previous century precisely to admit such generalizations.

1.10. Main Ideas of a Material Theory of Induction

The preceding exemplifies how I believe we should understand inductive inference. Let me collect the main ideas here:

Inductive inferences are warranted by facts not by formal schema.

What makes the inductive inference a good and strong one is not conformity with some universal formal schema; it is the facts pertaining to the subject matter of the induction. Hence, the warrant is “material” and not formal. Curie already knew of the closeness of the chemical properties of barium and radium. She knew of the well-established isomorphism that arose in such cases and indicated a closeness of the corresponding crystalline structures. Those facts assured her that the few cases she had observed of similarity between radium and barium chloride crystals could be generalized.

The essential idea here is that facts can serve a dual role, both as statements of fact and as warrants of inference. This idea is actually quite familiar. In deductive logic, the conditional “If *A* then *B*” serves this dual role. It can serve as a factual premise in an argument; or we can take the same argument and understand its role as warranting a deductive inference from *A* to *B*.

In chemistry, the facts that play this dual role look, loosely, like “Generally, *X*.” For example, “Generally, salts that are chemically analogous have similar crystalline structures.” This is both a fact in chemistry and an authorization to infer that radium salts and barium salts will have similar crystalline structures because of their chemical similarity. The inference is authorized all the more strongly when Curie found a single sample of radium chloride crystals that, as expected, exactly resembled barium chloride crystals. This diminished the possibility of smaller but superficially detectable differences. The inference is inductive since the chemical facts do not deductively entail Curie's inference. This is the import of the modifier "generally." It accommodates the ways the inference can still fail that are peculiar to this particular chemical example.

All induction is local. It is contextual.

The chemical facts that authorize these inductive inferences are truths of a particular domain of chemistry. They warrant a local mini-logic, peculiar to the context, in which evidence of chemical similarity and of a few samples warrants the generalizations indicated. This local mini-logic resembles the universal schema of enumerative induction. But the resemblance is superficial. There will, no doubt, be other domains in which other facts will warrant inferences that also resemble enumerative induction. The inferences of each domain will be distinct, carrying their own unique restrictions that do not derive from a universal schema, and bearing their own unique form of inductive risk.

Inductive inference is generically variegated and imprecise.

The imprecision here designates a lack of formal properties such as appear in mathematical theories of inductive inference. The inductive inferences on crystalline structure surveyed above can be characterized as "strong" or "reliable" or "very certain." These terms have a meaning only within the crystallographic context. Inferences to a unique crystallographic system are prone to failure if the salt displays polymorphism. The inference is "strong" just to the extent that polymorphism can be discounted.

Terms like these are variegated in that they have meanings peculiar to their contexts. The term "strong" will have one meaning in crystallography, another in some branch of physics, and yet another in some subfield of astronomy. What is missing generically is any precise means of comparing the strengths of inferences deemed "strong" in crystallography and in other domains, such as physics or astronomy. We also lack precise means of calibrating the difference between, say, "strong" and "very strong," within a single domain. This stands in contrast to contexts in which probabilities are applicable. The probability of at least one heads in ten coin tosses is 1/1,024 = 0.99902. In another domain, we may find that the probability that a parent passes on some specific genetic trait is 0.99. The two probabilities are comparable. The first exceeds the second by 1% and this slight difference will manifest eventually in slight frequency differences among many repeated trials.

The qualification "generically" allows that there are important exceptions. Background facts may sometimes authorize a precise, mathematical calculus of inductive inference. The most familiar case arises when we perform inductive inferences specifically on systems governed by probabilistic facts. Such systems include those undergoing radioactive decay, the forensics of DNA, and games of chance in a casino. Later chapters will describe systems in which other, non-probabilistic calculi of inductive inference are warranted. These precise calculi are only applicable when definite background facts warrant them.

The material theory does not authorize the default application of numbers to measure strengths of inductive support. It may be appealing to some to presume such numbers as a default. A probabilistic analysis can supply a definite number—say 0.99—whose closeness to unity gives the sought-for quantitative measure. As satisfying as this may be, without specific background facts to authorize the numbers, applying them is an exercise in spurious precision. It forces variegated notions of strength of support into a single, uniform mold that supposedly enables comparisons across domains. It neglects the domain-specific meaning for the strength of inductive support in each domain. To demand a single number or a single universal term to characterize inductive strengths across all domains invents a uniformity that is not found in the variegated character of inductive inference.

Inductive risk is assessed and controlled by factual investigation.

When one makes an inductive inference, one takes an inductive risk and one seeks both to assess and to minimize the risk taken. In a formal theory of induction, the assessment of the risk becomes an assessment of the reliability of the inference schema used. If we infer to the best explanation, we then need to ask how reliable it is to do that. And we are faced immediately with an intractable problem. There is no simple answer to this question; and there is likely no serviceable, complicated answer either.

In a material theory of induction, things are quite different. The warrant for an induction is a fact, and we assess and then control the inductive risk by exploring and developing that fact. Let us imagine that we notice that only a few radium chloride crystals resemble those of barium chloride. The inference to a broader resemblance might then be warranted by a chemical fact that salts manifest only a few crystalline forms. The strength of the inductive inference depends essentially on the correctness of that fact and just how many forms are admitted by the word "few." All of that can be checked by further investigation and just checking that is the normal business of research chemists. They developed theories of how crystals are constituted to enable a better understanding of which crystalline forms will appear in which circumstances. These investigations assure us that two salts will manifest similar crystalline forms if they are chemically similar; and this conclusion is in turn grounded both in other observations and a theoretical argument. Since radium and barium are chemically very similar, the chlorine atoms in a barium chloride crystal will permit the barium atoms to be replaced by radium atoms with minimum alteration to the crystal structure.

We assess and control inductive risk by learning more facts. The new facts provide new premises for inductive inference and new warranting facts. What was an intractable problem for a formal theory of induction becomes a routine problem in exploring the factual realm of chemistry for a material theory.

Inductive inference is material at all levels.

The crystallographic example explored here looks at particular sorts of inductive inferences at a specific level of refinement. One may wonder what happens if we take a more fine-grained view that looks more narrowly at specific inferences or—alternatively—if we take a coarser view that looks at inductive practice at a more general level. Will we find that a formal account of inductive inference succeeds there? Will we find that at levels of great refinement the glue that inductively binds the corpuscles of analysis is formal? Or will we find at a general level that a universal, formal theory emerges that can unify the diversity of the particular cases?

The claim here is that a material theory prevails at all levels. Of course, at all levels there will be inferences that loosely fit with one or other formal theory. We have seen in the case of crystallography that the inferences resemble enumerative induction. We should expect such loose fits, else the formal theories would not have survived in the literature. On closer examination, however, we will see that material facts are what warrant them.

1.11. Does the Material Theory Say That Inductive Inferences Are Really Deductive? No!

No. *No. NO.* It does not say that. This is perhaps the most frequent misreading of the material theory, and it can be put to rest here. The material theory maintains the distinction between the two forms of inference. In deductive inference, the truth of the premises assures the truth of the conclusion. In inductive inference, understood materially or otherwise, the premises only lend support to the conclusion. Inductive inference is not deductive inference.

The misreading of the material theory has it affirming that inductive inference is really some form of disguised deductive inference. My sense is that this misreading comes from a similarity between the material theory and another approach to inductive inference. In this other approach, we note that good inductive inferences are also deductive fallacies. For example, we take the following as a premise:

> This sample of salt *A* has crystallographic form *B*.

From this, we infer

> All samples of salt *A* have crystallographic form *B*.

This is a deductive fallacy. We could imagine that the argument is really, secretly a valid deductive argument, but we do not see it because one or more of the premises are unstated. That would make the argument an "enthymeme," a valid inference with unstated premises. In this case, a suitable unstated premise would be the strong form of Haüy's Principle:

Each crystalline substance has a single characteristic crystallographic form.

With this added premise, the inference becomes deductively valid. In the other approach, all inductive inferences are treated this way. They are treated as failed deductions that are repaired by supplying missing or unstated premises. This is not how the material theory treats inductive inference, however.

If we transform the inductive inference to a deductive inference by adding such premises, we have generated what is known as a "deduction from the phenomena." The best-known examples are given in Book 3 of Newton's *Principia*, where he shows how to infer deductively from the phenomena of celestial motions to the basic ideas of his theory of gravitation. His examples are so important that inferences of this type are often called "Newtonian deductions from the phenomena."

In admitting these cases, the material theory does allow that some inductive inferences may turn out to have been deductive inferences all along, once we make the background facts explicit.[12] However—and here is the key observation—this deductive outcome is an extreme and relatively rare case. Most commonly, it does not arise. When we identify the warranting facts, they supply an inductive warrant only. The strong form of Haüy's Principle is false. The correct, weakened form of Haüy's Principle merely asserts that "Generally, each crystalline substance has a single characteristic crystallographic form." The crucial word "generally" makes all the difference. It reminds us that the original principle fails if there is polymorphism. In accepting the conclusion, we take the risk that polymorphism—if present—will undo the conclusion. That is, the warrant supplied by the weakened form of the principle is not strong enough to assure us of the conclusion with deductive certainty. The distinction between deductive and inductive inference is maintained.

12 This is not a bad outcome at all. We thought that we must take an inductive risk in accepting the conclusion of the original inference. However, we learn that background facts assure us that no inductive risk is taken in accepting the conclusion. The inference has become deductive and, in effect, we already took the inductive risk needed when we accepted the background assumptions.

Chapters 2–9 will elaborate and illustrate these claims further through examination of a sequence of inductive inference forms employed in the literature: the replication of experiment, analogical inferences, inferences grounded in notions of simplicity, and inference to the best explanation. These chapters will be followed by an extensive investigation into the limitations of the Bayesian approach in Chapters 10–16. Where the present chapter has developed the material theory of induction by means of an example, the next chapter will develop the general arguments for it.

Note added March 15, 2020.
Commentaries on the draft chapters of this book have been collected for an issue of *Studies in History and Philosophy of Science*. It has become apparent from those commentaries that the draft chapters had not adequately distinguished two questions that arise within the material theory of induction. They are

(inductive-logical)
Question: Which inductive inferences are licit?
Answer: Those that are warranted by a (true) fact.

(epistemic)
Question: How can we know that a specific inductive inference is licit?
Answer: We must be assured of the truth of the appropriate warranting fact.

The first question is answered by matters of fact that obtain independently of any human beliefs, knowledge, or awareness. The answer to the second question depends on the answer to the first question. To know that some candidate inference is licit, we need to know the warranting fact. Gaining that knowledge can sometimes be troublesome. We may have to conjecture what the warranting fact is. In this case, we cannot be assured that the associated inference is licit until further investigation assures us that we have conjectured a factual truth.

REFERENCES

Accum, Frederick. 1813. *Elements of Crystallography After the Method of Haüy*. London: Longman, Hurst, Rees, Orme, and Brown.

Anon. 1911. *The Americana Supplement: A Comprehensive Record of the Latest Knowledge and Progress of the World Compiled by the Editorial Staff of the Americana Assisted by Expert Authorities Complete in Two Volumes*. The Scientific American Compiling Department.

Bird, Alexander and Emma Tobin. 2010. "Natural Kinds." *The Stanford Encyclopedia of Philosophy*. Summer 2010 Edition. Edited by Edward N. Zalta, http://plato.stanford.edu/archives/sum2010/entries/natural-kinds.

Boole, George. 1854. *An Investigation of the Laws of Thought*. London: Walton and Maberly.

Copi, Irving M. 1967. *Symbolic Logic*. 3rd ed. New York: MacMillan.

Curie, Marie. 1904. *Radioactive Substances*. 2nd ed. London: Chemical News Office.

———. (1911) 1999. "Radium and the New Concepts in Chemistry." In *Nobel Lectures: Chemistry 1901–1921*, pp. 202–12. Singapore: World Scientific.

Ford, William E. 1912. *Dana's Manual of Mineralogy.* 13th ed. New York: John Wiley & Sons.

Harman, Gilbert. 2002. "Internal Critique: A Logic is Not a Theory of Reasoning and a Theory of Reasoning is Not a Logic." In *Handbook of the Logic of Argument and Inference: The Turn Towards the Practical*, edited by R. H. Johnson, H. J. Ohlbach, Dov M. Gabbay, and John Woods, pp. 171–86. Amsterdam: Elsevier Science.

Haüy, René Just. 1807. *An Elementary Treatise on Natural Philosophy*. Vol. 1. Translated by Olinthus Gregory. London: George Kearsley.

Jeffrey, Richard C. 1983. *The Logic of Decision*. Chicago: University of Chicago Press.

Miers, Henry A. 1902. *Mineralogy: An Introduction to the Scientific Study of Minerals*. London: MacMillan.

Norton, John D. 2003. "A Material Theory of Induction." *Philosophy of Science* 70: pp. 647–70.

———. 2005. "A Little Survey of Induction." In *Scientific Evidence: Philosophical Theories and Applications*, edited by P. Achinstein, pp. 9–34. Baltimore: The Johns Hopkins University Press.

———. 2011. "Challenges to Bayesian Confirmation Theory." In *Handbook of the Philosophy of Science.* Vol. 7, *Philosophy of Statistics*, edited by Prasanta S. Bandyopadhyay and Malcolm R. Forster, pp. 391–439. Amsterdam: Elsevier.

Runge, C. and J. Precht. 1903. "The Position of Radium in the Periodic System According to Its Spectrum." *Philosophical Magazine* 5: pp. 476–81.

Rutherford, Ernest. 1913. *Radioactive Substances and their Radiations*. Cambridge: Cambridge University Press.

Whewell, William. 1837. *History of the Inductive Sciences.* Vol. 3. London: John W. Parker