

DRAFT

Chapter from a book, *The Material Theory of Induction*, now in preparation.

The Material Theory of Induction Stated and Illustrated

John D. Norton

Department of History and Philosophy of Science

University of Pittsburgh

<http://www.pitt.edu/~jdnorton>

The primary goal of this first part is to argue for a view of induction that I call the “material theory of induction.”¹ This first chapter will give a synopsis and illustrations of the theory. Later chapters will elaborate and support the view.

0. 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 at the outset what is meant by them here. Traditionally, induction has had a narrow meaning. At its narrowest, it refers to “induction by simple enumeration,” the inference from “Some As are B” to “All As are B.” This is an example of what is known as “ampliative inference,” for we have amplified the instances to which our knowledge applies. The premise applies just to the few cases of As at hand; the conclusion applies to all. I take this idea of amplification in its most general sense to be what induction is about. I shall use “induction” and “inductive inference” as the general term for any sort of ampliative inference; that is, any licit inferences that lead to conclusions logically stronger than the premises. Therefore the terms embrace what is sometimes called “abductive inference,” which is an inference to something that explains an otherwise puzzling phenomenon.

¹ For earlier accounts, see Norton (2003, 2005).

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, $P(H|E)$, 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 is changed. 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 7 to 5 to arrive at 12, we are merely 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. Under this alternative view, to say that we infer from proposition A to proposition B merely 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 when its normal usage is already well established. Could not another word have been found? Perhaps the redefinition is supported by the grammar of the use the term that presupposes an agent that infers. A similar redefinition might insist that saying “we add 7 to 5 to arrive at 12” merely reports our belief in the summation with no supposition that it conforms with arithmetic. I would find that redefinition equally troublesome.²

² Harman (2002, p. 173) gives a clear statement of the psychologized notion of inference that is *not* employed in this text: “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 long-standing and pervasive usage of “rules of inference” as designating licit manipulations and argument schemas, such a modus ponens and various syllogisms. See, for example, Boole (1854, Ch. XV) and Copi (1967, p.36 and inside back cover).

Throughout this volume, unless some context demands a momentary 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. 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 72 eggs overall. It also tells us that if we have six troupes of a dozen acrobats, then we have 72 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

$$6 \times 12 = 72$$

That is really a schema that says (amongst other things)

If you have 6 *groupings* of 12 *individuals*, then you have 72 *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 don’t come with restrictions to particular domains. It is the same arithmetic for eggs and acrobats. And the rules are formal in the sense that they attend only to the form of the sentence asserting the data: six of 12

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.

2. Problems of the Formal Approach

However it 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 an explanation better or even what an explanation is. Efforts to make these notions precise open more problems than they solve. Or we are told that all of inductive logic is subsumed by probability theory. A later part of this book is 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 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? A formal theory of induction distinguishes the good inductive inferences from the bad by means of universal schemas. In its place, I urge a material theory of induction. According to it, what separates the good from the bad inductive inferences are background facts, the *matter* of the inference of the inference, as opposed to its *form*. Or, 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.

3. Inductions on Crystal Forms

An example will make the problems of the formal accounts 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. Let us say that a chemist prepares a new salt of some metal and notes its particular crystalline form. It is routine for the chemist to report the form not merely as the form of this sample, but as the form of this 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 it is easy to overlook its inductive character, we should expect a good treatment of it from an account of inductive inference.

To develop the example, we need to appreciate that 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 hard to find a simple and robust system of classification. That complication will prove to be important for inductive inference.

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 three perpendicular axes of equal length. The same is true of a regular octahedron, which also conforms to the system. There are many more shapes in the system. Crystals have natural cleavage planes. If we start with a crystal cube of common salt and begin chipping off the crystal's corners along these planes, we can chisel it

down to a regular octahedron. The cube on the left of Figure 1 has had its corners chiseled off, on the way to the octahedron shown inside.




Figure 1. Chiseling Off the Corners of a Cube³

During the process of chiseling, we will pass through many, more complicated shapes of cubes with corners removed to different extents. The shape on the right of Figure 1 is such an intermediate form. All these multi-faceted shapes and many more are licit forms within the cubic system. Figure 2 shows further crystal shapes, both within the cubic system. They derive from the octahedron shown within to which they may be chiseled down.




Figure 2. Mis-shapen Octahedra⁴

³ Redrawn after Miers (1902, p. 14, Fig. 17 and 18).

⁴ Redrawn after Miers (1902, p. 11, Fig. 9 and 10).

All these shapes are different from the crystalline shapes permitted to barium chloride, for barium chloride is monoclinic. That means that its crystals are characterized by three unequal axes, two of which intersect at an oblique angle and the third 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 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 3, the inclination is towards the right of the figure.




Figure 3. Primitive Form of the Monoclinic System

The range of crystal shapes allowed in the monoclinic system are derived from this form, in the same way that those allowed in the cubic system are derived from 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.

4. Curie and Radium

Perhaps the most famous of all episodes in crystal formation was Marie Curie's separation of radium by fractional crystallization from uranium ore. 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. 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 it as a trace impurity, revealed by its great radioactivity.

The final separation of the radium chloride from the barium chloride is 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, those crystals will harbor more radium chloride. The solution remaining above the crystals has a fifth the radioactivity of the original, Curie reported. While that seems like a large increase, the quantity of radium present is so tiny that it falls far short of what is 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 this is described in her doctoral dissertation (Curie, 1904), presented to the Faculté des Science 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 (Curie, 1904, p. 26) and they bleed off into less certain reports on the colors of the crystals that, she suspects, may 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 and, soon, others separated out only miniscule 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, Rutherford (1913, p. 470) 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.

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

5. 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 is the formal analysis. 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 for the As and Bs would yield a feeble induction. To get an induction of the strength seen by Curie and Rutherford, we have to be very selective in what is substituted for A and B. The As 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 very carefully chosen properties B. There are very many ways of describing crystal forms. Virtually none of them support a strong inductive inference.

To revert to the simpler example, one may find 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 that property can be inserted into the schema of enumerative induction to form the generalization.

This 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 published in the mineralogical literature. His *History of the Inductive Sciences* (1837, Vol. III, Book XV, Ch.1-2) gives a lively account of these hesitations and their clarifications by Romé de Lisle and René Just Haüy after 1780.

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 above 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.

Isomorphism is a term of art then and now used to describe the circumstance in which two different substances have very close chemical and crystalline properties. (See, Miers, 1902, p. 213.) It saved Curie 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 tells us that radium chloride has it too.

If the schema of enumerative induction is to function as a general logic, these 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:

“...and, 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. That is not a logic, but merely a catalog whose guiding rationale is 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 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 is a natural kind term. The familiar characterizations of natural kind terms include that the terms support induction. (Bird and Tobin, Section 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 to say in fancy words 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?

6. A Bayesian Attempt

These sorts of problems will defeat other attempts to provide a formal account of the inference. Here is a quite different attempt at another extreme of the literature.

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 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 $E = E_1 \& E_2 \& \dots \& E_n$, where E_i asserts the evidence in the i -th case. The probability of interest is $P(H|E)$, the probability of the hypothesis H given the evidence E . It represents the inductive support afforded to H by E , if we think of the probabilities objectively. Or it is the belief we have in H given that we know E , if we interpret the probabilities subjectively. We are interested in seeing how the posterior probability $P(H|E)$ compares with the prior probability, $P(H)$; that is, we seek how is the probability of H changes when we incorporate our learning of evidence E . Those 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

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

We will not be able to compute $P(H|E)$ 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 prior probabilities, $P(H)/P(\sim H)$ changes to $P(H|E)/P(\sim H|E) = r$. From this last ratio, $P(H|E)$ can be recovered as

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

The theorem tells us that the controlling quantities are the two likelihoods, $P(E|H)$ and $P(E|\sim H)$. 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 : $P(E|H) = 1$.

The other likelihood $P(E|\sim H)$ is much harder to determine. How probable, it asks us, is the evidence if the hypothesis is false? Answering that requires some creative imagination for we have no precise prescription for the ways that 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 $P(E|\sim H)=0$. Inserting this into Bayes' theorem leads to $P(H|E) = 1$; the hypothesis is maximally probable.

However things are not quite so simple. 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. That means that the salt can crystallize into two or more the systems. So there is some chance, perhaps small, that the evidence E_i obtains, even if H is false. Let us set that probability to q so that $P(E_i|\sim H)=q$. Let us suppose that each of the samples is taken under independent conditions, so that the obtaining of each E_i is probabilistically independent of the others and the probability of the conjunction is a just a simple product of terms:

$$P(E|\sim H) = P(E_1 \& E_2 \& \dots \& E_n|\sim H) = P(E_1|\sim H).P(E_2|\sim H). \dots .P(E_n|\sim H) = q^n$$

Bayes' theorem now becomes

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

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 $r = P(H|E)/P(\sim H|E)$ also grows arbitrarily large. That corresponds to the posterior, $P(H|E) = r/(r+1)$ coming arbitrarily close to unity. And that 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. That is for the better, since the prior probabilities $P(H)$ and $P(\sim H)$ would have to be plucked from the air.

7. What is Wrong With It

If one inclines to numerical and algebraic thinking, this may seem like a very satisfactory analysis. It has brought mathematical precision to what 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. However the harm done is great, since we have convinced ourselves that we have solved a great problem, when we have not.

First, the analysis is troubled by problems on the technical level. The calculation can only succeed if we make the right sorts of assumptions. We have to assume that each of the observations E_i are probabilistically independent given that H is false. Need I say that speculating on how things might be connected if the hypothesis is false is a risky operation, for hypotheses can fail in many ways? Setting that aside, in the case of radium salts, the samples are all recovered in similar circumstances, by fractional crystallization of traces from a solution of the corresponding barium salt. Given that similarity in each case, there is every reason to expect that the observations are *not* independent. What if radium salts turn out to be polymorphous so that they may crystallize in several systems? The common conditions of radium preparation may favor just one, so that we never see the other systems realized. That would be a failure of independence.

Second, 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 see if 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: it introduces specific probability numbers, while no such numbers are in evidence in the inductive practice. Just which value is appropriate for $P(E_i|\sim H)$? 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 are merely reflections of 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, exemplar computations 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 head (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 independence is doubtful, then repeat the analysis assuming some form of dependence between the evidence instances E_i . Or if selecting appropriate prior probabilities is troublesome, then devise rules or systems for selecting them. This maneuver does not solve the problem. It merely enlarges the analysis and exiles the same problems to remote corners, where the same problems remain and more appear. They are just harder to see because the analysis has become so enlarged and so much more complicated.

The third problem is, in my view, the most serious. It cannot be deflected by the strategy of expand and exile. 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 that description is given, the most important part of the inductive analysis is over. 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 describe things in these terms, no analysis will work, including 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 so that no single description fits and so that we do not even have 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 was 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 the 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 work lies in finding the appropriate system of classification. In introducing crystallography as a “classificatory science,” Whewell (1837, pp. 212-13, his emphasis) stresses 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*.

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

8. 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. Examples are conforming 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.

Chemists at the start of the 20th 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 were laid by René Just Haüy in the late 18th and early 19th 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. Common salt, for example, is built from minute cubes. The many shapes of crystals of common salt were just those shapes that can be constructed by stacking up these little cubes. They include large cubes and anything that can be derived from it by fracturing pieces off. The oblique faces admitted by fracturing off the corners of the cube are, at the smallest scale, really many staircases of these cubes. But that scale is so small that we perceive a perfectly smooth surface. In his treatise published at the time Curie was working on radium, Henry Miers (1902, p.48) illustrated Haüy's account as in Figure 4:⁵




Figure 4. Haüy's Account of Crystal Fracture Planes

An account⁶ contemporary to Haüy summarized the theory (Accum, 1813, p. 110):

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.

⁵ The figure on the left is redrawn from Mier's Fig. 38 and the figure on the right is a reproduction of Mier's Fig. 37.

⁶ This account is more succinct than Haüy's own synopsis of Haüy (1807, pp. 86-101).

From this theory came the essential result that every substance was characterized by a unique primitive form (Accum, p. 117)

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.

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. Accum (1813, p. liv) listed ten substances based on the cube, for example. Among them are 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 as (Anon, p. 365 under "crystallography):

The Abbé Reny Just Hauy, 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.

The view outlined was not so much a principle as a simple consequence of his theory; and Haüy's theory, as outlined by 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 is sufficient to identify the crystallographic system to which all crystals of that substance must belong. We have the inference:

(Haüy's Principle) Each crystalline substance has a single characteristic crystallographic form.

This 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 no 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 passage was possible only during a brief window of a few decades of 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 (837, p. 235) reports 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 the one crystalline substance may for crystals belonging to two, three or many crystallographic systems—called “dimorphism,” “trimorphism” and “polymorphism” respectively. It was not clear how merely stacking the nuclei of the same shape could yield these different shapes. Mineralogy texts of the early 20th century routinely reported examples. Ford's (1912, p. 80) list is presented more as a reminder of what everyone supposedly knew, 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.

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.

However 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 as

(Weakened Haüy's Principle) *Generally*, each crystalline substance has a single characteristic crystallographic form.

This 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 insertion of 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. Perhaps its schema is:

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.⁷

This is one illustration of how background knowledge drives inductive inferences and how that 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 at the time Curie worked the isomorphism, mentioned earlier. This was then defined more precisely as (Ford, 1912, p. 79):

A series of compounds which have analogous chemical compositions and closely similar crystal forms are said to make an isomorphous group.

An early celebrated instance was a triumph of crystallographic analysis. Whewell (1937, pp. 226-28) reports confusion over the crystalline substance “heavy spar.” Haüy found that its cleavage angles varied by three and a half degrees, according to the origin of sample. One was from Sicily and one from Derbyshire. That 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

⁷ While the inferences may look formally similar, they will be quite different if applied to crystals or to astronomy. Take the proposition: Generally, orbiting objects in our solar system orbit in the same direction as the earth. From it, we may infer with a small risk, that this recently discovered asteroid will orbit in that same direction. The risk we are taking is distinct from that taken in crystallography. We are risking the possibility that this asteroid was not formed by the same processes that formed most other objects in our solar system.

out that the two samples were of different substances. The Sicilian was barium sulphate; and the one from Derbyshire strontium sulfate. 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, as was essential to this story, not perfectly identical. They are a classic case of isomorphism.

When Curie remarked that the radium chloride formed crystals having “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 hard. Hence the familiar idea of isomorphism would have indicated that the crystals of the two chlorides should be similar. All that was really left to affirm was how close the similarity would be. It was, Curie found, “exactly the same.”

Immediately after Curie’s work, 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. (Soddy, 1907, p.332) 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

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 was 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 a “some” to an “all” is so strong as to be essentially unquestioned. On closer inspection we found that appearance quite 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 it within a complicated network of chemical ideas that had been devised over the previous century precisely to admit such generalizations.

9. Main Ideas of a Material Theory of Induction

This last case exemplifies how I believe we should understand inductive inference. Let me collection the main ideas:

Inductive inferences are warranted by facts not formal schema.

What makes the inductive inference a good and strong one is not conformity with some universal formal schema. It is 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 of 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 warrants of inference. That idea is actually quite familiar. In deductive logic, the conditional “If A then B.” Serves that 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 resembles barium chloride crystals. That diminished the possibility of smaller but superficially detectible differences. The inference is inductive since the chemical facts do not deductively entail Curie's inference. That is the import of the "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 inference 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 are 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 another in some subfield of astronomy. What is missing generically is any precise means of comparing the strengths of inferences deemed as "strong" in crystallography and in some disjoint domain, such as physics or astronomy. We also lack precise means for calibrating the difference between, say, "strong" and "very strong" within a single domain. This stands in contrast with contexts in which probabilities are applicable. The probability of at least one head in ten coin

tosses is $1/1024 = 0.99902$. 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, or the forensics of DNA or 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 it may be, without specific background facts to authorize them, applying these numbers 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, that 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. We are faced immediately with an intractable problem. There is no simple answer to this question; and 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 the fact. Let us imagine that we notice only that 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 “few.” All that can be checked by further investigation and just 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 in both 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 crystal structure.

We assess and control inductive risk by learning more facts. These new facts both provide new premises for inductive inference and also new warranting facts. What was an intractable problem for a formal theory of induction has become 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 even more narrowly at very specific inferences; or if we take a coarser view that looks at inductive practice at a more general level. Might we find a formal account of inductive inference succeeding there? Might we find that, at levels of great refinement, the glue that inductively binds the corpuscles of analysis is formal? Or that, at a very general level, a universal, formal theory emerges that can unify the diversity of the particular cases.

The claim 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 crystallographic case that the inferences resemble enumerative induction. We should expect such loose fits, else the formal theories could not have survived at all in the literature. However, they will always be loose fits and, I maintain, closer examination will reveal that material facts are warranting them.

10. Does The Material Theory Say that Inductive Inference is 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 as a premise:

This sample of salt A has crystallographic form B.

and from it 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 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 this other approach, all inductive inference is treated this way. They are treated as failed deductions that are repaired by supplying missing or unstated premises. It 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 III 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 deduction 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.⁸ 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 Häüy’s Principle is false. The correct weakened form of Häüy’s Principle merely asserts that:

Generally, each crystalline substance has a single characteristic crystallographic form. That 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 after all, 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.

---oOo---

The chapters to come 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. They will be followed by an extensive investigation into the limitations of the Bayesian approach. 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.

References

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

⁸ That 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 have already taken any needed inductive risk when we accepted the background assumptions.

- 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 Tobin, Emma, (2010) "Natural Kinds", *The Stanford Encyclopedia of Philosophy* (Summer 2010 Edition), Edward N. Zalta (ed.), <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.
- Ford, William E. (1912) *Dana's Manual of Mineralogy*. 13th ed. New York: John Wiley & sons.
- Haüy, René Just (1807), *An Elementary Treatise on Natural Philosophy*. Vol. 1. Trans. Olinthus Gregory. London: George Kearsley.
- Harman, Gilbert (2002) "Internal Critique: A Logic is not a Theory of Reasoning and a Theory of Reasoning is not a Logic," pp. 171-86, in, R.H. Johnson, H.J. Ohlbach, Dov M. Gabbay, John Woods, eds., *Handbook of the Logic of Argument and Inference: The Turn Towards the Practical*. Amsterdam: Elsevier Science.
- 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.
- Norton, John D. (2005) "A Little Survey of Induction," in P. Achinstein, ed., *Scientific Evidence: Philosophical Theories and Applications*. Johns Hopkins University Press, 2005. pp. 9-34.
- Runge, C and Precht, J, (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.
- Skłodowska Curie, Marie (1904) *Radioactive Substances*. 2nd ed. London: Chemical News Office.
- Skłodowska Curie, Marie (1911) "Radium and the New Concepts in Chemistry," pp. 202-212 in *Nobel Lectures: Chemistry 1901-1921*. Singapore: World Scientific, 1999.
- Whewell, William (1837) *History of the Inductive Sciences. Vol. III*. London: John W. Parker