

Induction without Probabilities¹²

John D. Norton

Center for Philosophy of Science

and Department of History and Philosophy of Science

University of Pittsburgh

www.pitt.edu/~jdnorton

A simple indeterministic system is displayed and it is urged that we cannot responsibly infer inductively over it if we presume that the probability calculus is the appropriate logic of induction. The example illustrates the general thesis of a material theory of induction, that the logic appropriate to a particular domain is determined by the facts that prevail there.

1. Introduction

What is the right logic of induction? In recent decades, one answer has taken the first position in philosophy of science. It is that the probability calculus is *the* logic of induction. Indeed the success of this idea has been so marked that many feel it is the only game in town. This sentiment was already encoded over a decade ago in the 1992 title of John Earman's *Bayes or Bust*. My purpose in this note is to argue that this view is entirely too narrow. While the probability calculus is the appropriate logic of induction in many important cases, it is not the universal logic of induction.

My argument will begin with a challenge in Section 2 below. I will display a system in which the reader is invited to infer inductively. My claim will be that this cannot be done

¹ I am grateful to Fred Kronz for formal comments on an earlier version of this text at the conference "The War on Induction," Johns Hopkins Center for History and Philosophy of Science, March 23, 2007.

² For an elaboration of the views sketched here, see Norton (2003, 2005, forthcoming).

responsibly if the inductive logic must be probabilistic. For readers impatient to know the trick, it is simple. I will invite readers to infer inductively over an indeterministic system for which the full specification of the present state does not enable the laws of nature to assign physical chances to the many possible futures. Assigning probabilities to those futures requires us to know more than Nature says we can know.

In the sections immediately following, I will lay out the background for this challenge. In Section 3, I will sketch out why I am unmoved by the many arguments in the literature, each purporting in its own way to establish that the probability calculus is the One True Logic of Induction. In Section 4, I will review why I believe that a logic of induction must be responsive to the physical facts governing systems over which we infer. Lewis' "principal principle" is an early form of that idea, which I have developed more extensively in what I call a "material theory of induction." It asserts that the logic of induction appropriate to some domain is determined by the material facts that prevail in that domain. It follows that there is no single logic of induction applicable in all domains. That approach will then be used in Section 5 and 6 to display what I take to be a responsible way to infer inductively over the system described in Section 2. In Section 7, I will display a precise sense in which the belief distribution generated in Section 6 corresponds to a state of complete ignorance. To do so, I will recall recent work in which I argued that we have two means, each able to identify this state already in the literature. One is the principle of indifference; the other is invariance under symmetric redescription. Section 8 will offer concluding remarks.

2. A Puzzle for Inductive Inference

In most branches of physics, we can find systems, usually highly idealized, that are indeterministic. They are systems for which a full specification of their present state fails to fix what their future state will be. We have long been familiar with one manifestation of indeterminism. In quantum theory, in most cases, a specification of the present state of the system only fixes the probabilities of different futures. In a more extreme form of indeterminism, the full specification of the present leaves the future undetermined and—the key fact of importance here—our physical theories provide *no physical chances* for the different futures. They tell us only which futures are possible. Some recent examples arise in the supertask literature. See, for example, Alper et al. (2000); Norton, (1999). One of the simplest examples in

Newtonian physics is “the dome” and it will be my focus here. See Norton (2003a, §3) for a more complete description; and Norton (manuscript a) for a discussion of why the dome ought to be characterized as properly a part of Newtonian physics.

While nothing in my argument depends upon this example specifically, I will develop my story using it alone just for concreteness.³ Otherwise, I do not intend my moral to be restricted just to inductive inference in the physical sciences. Rather I intend it to have universal application. The reason I develop an example in the physical sciences is that in it there is a minimal level of vagueness over the details of the relevant facts in the world. In other cases in which these facts are less clearly articulated, that very vagueness can make determination of the appropriateness of different sorts of inductive inferences very difficult.⁴

A point mass can slide frictionlessly over a dome with circular symmetry in a vertical gravitational field. Initially, the mass is motionless at the apex. See Figure 1. If the shape of the surface is chosen appropriately, Newton’s equations admit many solutions. The mass may remain at rest indefinitely at the apex; or it may remain at rest for some arbitrary time T and then spontaneously accelerate in any radial direction.⁵ It is important to note that the spontaneous motion does not arise from some very slight perturbation, a miniscule wobble, say, that shakes the mass free at the moment of spontaneous excitation, time $t=T$. Nothing changes in conditions of the dome. It is just that Newton’s equations of motion admit multiple solutions, one in which the mass remains at rest at times $t>T$ and one in which it moves for $t>T$.

³ Nothing in my argument requires the reader to adopt a strong view over determinism or indeterminism in physics. All it needs the reader to accept is that some scenario with non-probabilistic indeterminism is at least conceivable and that we would expect inductive inference to be applicable in that scenario.

⁴ An illustration: A once followed B. Should we now infer that an A always follows a B?

⁵ This happens if the surface is such that $h = (2/3g)r^{3/2}$, where r is the radial distance in the surface of the dome and h the vertical distance below the apex; g is the acceleration due to gravity. For a unit mass on the surface, Newton’s laws entail an outwardly directed acceleration field, $F = d(gh)/dr = r^{1/2} = d^2r/dt^2$. This equation is solved by $r(t)=0$, for all t ; and by a spontaneous excitation at T : $r(t) = 0$, for $t \leq T$ and $r(t) = (1/144)(t-T)^4$, for $t \geq T$.




Figure 1. The dome: an indeterministic system⁶

How are we to infer inductively over the time T of spontaneous acceleration? If we are given evidence E that the mass is at rest at time $t=0$, how much inductive support does E accord to hypotheses $H(T_1, T_2)$, that the mass will begin to move at the time t in the interval $T_1 \leq t < T_2$?

3. Failure of Arguments to Prove Universality of the Probability Calculus as the Unique Logic of Induction⁷

There have been numerous attempts to establish that the probability calculus is the universally applicable logic of induction. The best known are the Dutch book arguments, developed most effectively by de Finetti (1937), or those that recover probabilistic beliefs from natural presumptions about our preferences (Savage, 1972). Others proceed from natural supposition over how relations of inductive support must be, such as Jaynes (2003, Ch. 2).

3.1 Working Backwards

These demonstrations are ingenious and generally quite successful, in the sense that accepting their premises leads inexorably to the conclusion that probability theory governs inductive inference. That, of course, is just the problem. The conclusion is established only in so far as we accept the premises. Since the conclusion makes a strong, contingent claim about our world, the demonstrations can only succeed if their premises are at least strong factually.⁸ That

⁶ Figure from Norton (2003a, §3).

⁷ Material in this section is drawn directly from Norton (forthcoming, Section 2).

⁸ There is no escape in declaring that good inductive inferences are, *by definition*, those governed by the probability calculus. For any such definition must conform with essentially the same facts

makes them at least as fragile as the conclusion they seek to establish. Since they are usually created by the simple expedient of working backwards from the conclusion, they are often accepted just because we tacitly already believe the conclusion.

For these reasons, all demonstrations of universality are fragile and defeated by a denial of one or more of the premises. A few examples illustrate this general strategy for defeating the demonstrations. Dutch books arguments are defeated simply by denying that some beliefs are manifested in dispositions to accept wagers. Or their results can be altered merely by adjusting the premises we will accept. Dutch book arguments commonly assume that there are wagers for which we are willing to accept either side. That assumption is responsible for the additivity of the degrees of belief the argument delivers. Its denial involves no incoherence, in the ordinary sense. It just leads us to a calculus that is not additive. (See Smith, 1961.) Similarly, there is no logical inconsistency in harboring intransitive preferences. They will, however, not sustain a recovery of transitivity of beliefs in Savage's (1972, §3.2) framework, which is necessary for beliefs to be probabilistic.⁹ Finally, Jaynes (2003, §2.1) proceeds from the assumption that the plausibility of A and B conditioned on C (written "(A|C)") must be a function of (B|C) and (A|B|C) alone, from which he recovers the familiar product rule for probabilities, $P(A|C)=P(A|B|C)P(B|C)$. That this sort of functional relation must exist among plausibilities, let alone this specific one, is likely to be uncontroversial only for someone who already believes that plausibilities are probabilities and has tacitly in mind that we must eventually recover the product rule.¹⁰

in that it must cohere with canonical inductive practice. Otherwise we would be free to stipulate any system we choose as the correct logic of inductive inference.

⁹ Savage's framework harbors a circularity. In its barest form, it offers you a prize of \$1, say, for each of the three acts f_A, f_B, f_C , if uncertain outcomes A, B or C happen, respectively. You prefer f_A to f_B just in case you think A more likely than B. So your preferences on f_A, f_B and f_C will be transitive just in case you already have transitive beliefs on the possibilities of A, B and C.

¹⁰ A simple illustration of an assignment of plausibilities that violates the functional dependence is "Plaus." It is generated by a probability measure P over propositions A, B, ... as a coarsening,

The fragility of these demonstrations is very similar to the failure of attempts to derive Euclid's fifth postulate of the parallels from the other four postulates. They all depended on displaying a logically equivalent assertion in geometry, whose denial, we are to suppose, is in some sense incoherent. It was eventually realized in the nineteenth century that the denial of these equivalent assertions involved no inconsistency; it merely led us to different geometries.

While I believe all these demonstrations fail in establishing universality, they still have great value. For we learn from them that, in domains in which their premises hold, our inductive inferences must be governed by the probability calculus.

3.2 The Surface Logic.

There is a second sort of argument for universality, mostly suggested indirectly by impressive catalogs of the success of Bayesian analysis at capturing our intuitions about inductive inference. All these intuitions so far have been captured by the probability calculus; so, the thought goes, we should expect this success to continue.

In my view, the success is overrated and does not sustain the probability calculus as the unique logic of induction. In many cases, the success is achieved only by presuming enough extra hidden structures—priors, likelihoods, new variables, new spaces—until the desired intuition emerges. That does not mean that the logic on the surface is probabilistic, but only that this surface logic can be simulated with a more complicated, hidden structure that employs probability measures.

Two examples will illustrate the concern. Take Hempel's original question of whether a non-black, non-raven confirms that all ravens are black. A probabilistic analysis gives an intuitively very comfortable result. But it only succeeds by adding a great deal of new structure to the original problem: populations with different distributions of ravens and black objects and a presumption that we are sampling randomly from them. That changes the problem to a new one amenable to probabilistic analysis. (For a survey, see Earman, 1992, §3.3) Consider ignorance, which, I argue below in Section 7, is not represented in an additive calculus. It may be introduced by associating beliefs with convex sets of probability measures. While additive

with only two intermediate values: $\text{Plaus}(A|B) = \text{"Low"}$ when $0 < P(A|B) < 1/2$; and $\text{Plaus}(A|B) = \text{"High"}$ when $1/2 \leq P(A|B) < 1$

measures were used to produce them, the sets themselves no longer conform to a logic with the algebraic property of *Addition* as defined below. Additive measures are merely the device used to generate a new system governed by a different surface logic.

Once again there is a geometric analogy. We can recover many non-Euclidean geometries by considering curved surfaces embedded in a higher dimensioned Euclidean space. That does not mean that Euclidean geometry is the universal geometry. It is not the geometry intrinsic to the surface. However we learn that Euclidean geometry can be used as a tool to generate that geometry, as could other geometries.

4. Material theory of Induction

How should we infer inductively over the dome? We can take inspiration from an idea introduced by David Lewis (1980) and now widely accepted. According to his “principal principle,” if we are inferring inductively over systems for which physical chances are available, we should conform our degrees of belief to those physical chances. A familiar example of such a system is an atom of a radioactive element undergoing radioactive decay. It is governed by the law that the physical probability $P(t)$ of the atom decaying in time t is

$$P(t) = 1 - \exp(-t/\tau). \quad (1)$$

where the time constant τ is related to the element’s half life by $t_{1/2} = \tau \ln 2$. So we should conform our degrees of belief to that formula. The magnitude $P(t)$ should be our degree of belief in the atom decaying within time t .

What of a case in which no physical chances are to be had? What if the system is indeterministic, but the laws governing the multiplicity of possible futures are non-probabilistic and assign no physical chances to different possible futures? The essential insight that lies behind Lewis’ “principal principle” still applies. If the realm of possibility is fully governed by some natural law, we should allow that law to dictate the course of our inductive inferences. Since this is just the case of the dome, in the next section, I will use the structure of possibility that Newtonian theory supplies as the framework for the inductive logic of the dome.

This last idea is an illustration of the general approach that I have taken elsewhere to inductive inference. According to a material theory of induction (Norton, 2003, 2005), the logic of induction governing each domain is fixed by the material facts that obtain in that domain. Two consequences follow immediately concerning inductive inference. First, since the pertinent

material facts vary from domain to domain, there is no universal logic of induction. Each domain has its own inductive logic picked out by the material facts that prevail there. In a slogan: all induction is local. Second, the demonstration of the validity of an inductive inference is not the displaying of a universally applicable inductive inference schema to which the inference at hand conforms. Rather the process terminates in the identification of material facts that assert the admissibility of the inference.

This domain dependence of inductive inference can be illustrated by inductive inferences on electrons. We routinely conduct a very secure inductive inference from the (rest) mass of one, or a few, electrons to the mass of all electrons. Thompson's original identification of the mass to charge ratio of the electron and Millikan's identification of the electron's mass and charge depended on determining those values in just a few cases and then, without apology, generalizing from them to all electrons. The usual explication of this inference is that it is an instantiation of enumerative induction, the inference from "some..." to "all..."; and one of the few reliable ones. A material theory of induction denies that enumerative induction is a universally applicable schema. It allows, however, that it is applicable in this domain for this property. The reason is that it is licensed in this particular application by the material fact that all electrons are fundamental particles and, generally speaking, all fundamental particles of the same type have the same mass.¹¹

The corresponding inference is not licensed for the momenta of electrons. What can be inferred from knowing that one or more electrons have such-and-such momentum depends very much on the facts prevailing in the relevant domain. If, for example, these are electrons that are part of a thermal system, then their momenta will be distributed according to the probabilistic laws of statistical physics. Those probabilistic laws could then serve as material facts that license a probabilistic analysis analogous to radioactive decay sketched above.

¹¹ A common response to examples like this is that the material theory turns all inductive inference into deduction. The essential qualification in this material fact is that *generally speaking* all fundamental particles of the same type have the same mass. This sameness of mass turned out recently to fail for neutrinos that come in varieties with slightly different masses. As a result, the material fact can only make the conclusion very likely, so that some inductive risk is taken in accepting it.

5. How NOT to Infer Inductively about Indeterministic Systems

Consider once again the decay of an atom of a radioactive element. What support does the evidence E that the atom is undecayed at time $t=0$ give to the hypothesis H(T) that the atom decays over the following time $0 \leq t < t$? Henceforth, let us write that degree of support as

$$[H(0, T)|E]$$

Since the process is fully governed by the law of radioactive decay (1), the material theory of induction enjoins us to set the degree of support equal to physical probability of decay

$$[H(0, T)|E] = P(T) = (1 - \exp(-T/\tau))$$

so that these degrees of support conform to the probability calculus.

Now consider the inductive inference problem concerning the indeterministic dome posed in Section 2. We have

E: At $t=0$, the mass is motionless at the apex.

H(T_1, T_2): The mass will begin to move at the time t in the interval [T_1, T_2), that is, in $T_1 \leq t < T_2$?

What is the degree of inductive support $[H(T_1, T_2)|E]$ accorded by the evidence E to the hypothesis H(T_1, T_2)?

Many find in the analysis of the radioactive decay of an atom a template that they cannot resist applying to the dome. They propose

$$[H(0, T)|E] = P(T) = (1 - \exp(-T/\tau))$$

The motivation is that the law of radioactive decay has an important property. It is the *unique* decay law that has a “no memory” property. If the atom has not decayed after 1 time unit, or 5 time units, or 100 time units, then the probability of decay in the next unit of time is still the same. Speaking metaphorically, it is as if the atom does not remember how long it has survived without decay, when it decides whether to decay in each new unit of time.¹²

¹² The property is seen most easily by considering $Q(t) = 1 - P(t) = \exp(-t/\tau)$, the probability of *no* decay in an initial time t . The probability of no decay in time u subsequent to a period t of no decay is just

A distinctive feature of the dome is that it also has this "no memory" property. Whether the spontaneous motion happens at some moment is quite independent of how long the mass has been sitting at the apex. So, if any probabilistic law is applicable to the dome, it is this one. However we cannot set our degrees of support [H|E] equal to probabilities governed by the same formula as in the law of radioactive decay. Any instance of the law of radioactive decay has a time constant τ in it. That time constant exercises a powerful influence on the chances of the spontaneous event. Figures 2, 3 and 4 display graphs of $P(t)=P(H(0,t)|E)=P(H|E)$ for values of $\tau=0.1$, $\tau=1$ and $\tau=10$:




Figure 2. Decay with time constant 0.1




Figure 3. Decay with time constant 1

$$Q(t+u)/Q(t) = \exp(-(t+u)/\tau) / \exp(-t/\tau) = \exp(-u/\tau) = Q(u),$$

which is just the probability of no decay in an initial time u .




Figure 4. Decay with time constant 10

A very small time constant makes the decay very probable, virtually immediately; a very large time constant delays the decay very probably, for a long time.¹³ Nothing in the physics of the dome fixes a time constant or any sort of time scale for the spontaneous motion of the mass. The physics is completely silent on how soon the motion may happen. It just says "it's possible."

So if we are to use the probabilistic formula, we must add a time scale. That is, we must pretend to know more than the full physical specification of the problem allows. Speaking metaphorically, Nature, in the guise of Newton's physics, is unable to assign a time scale to the decay. If we assign one, we must pretend to know more than Nature. Our original goal was merely to reason inductively about a system. Yet we have ended up as physicists, proposing new physical properties that the system--by construction--does not have.¹⁴

There is a loophole. The probability formula (1) of the law of radioactive decay is the unique rule with the "no memory" property. So if any probability formula would work for the

¹³ Setting $\tau=\infty$ gives us the case of $P(T)=0$ for all T , so that the mass never moves. That is we assign unit probability to the outcome of no motion ever.

¹⁴ Might a probabilistic analysis be possible if only we were given a little more data, such as the results of observation of several domes over some time? This strategy only makes sense if one does not accept the initial supposition that the Newtonian analysis gives the full physics of the dome. If one accepts that it does, then no catalog of outcomes will give any new, useful information for the inference problem. The situation is analogous to someone keeping detailed records of the outcomes of a roulette wheel's spins in the hope that some dependence between successive spins will be manifested. That strategy in a casino is futile if the wheel has been properly constructed so that there is no dependence.

dome, that one would have to be the one. However statisticians sometimes use improper probability distributions—that is, ones that do not normalize to unity—and there is an improper density with this “no memory” property. It is the uniform density $p(H(0,t)|E) = dP(t)/dt$, shown in Figure 5, that assigns that the same small probability ϵ to each unit time interval.




Figure 5. An Improper Probability Distribution

It is "improper" since the probability assigned to all the unit time intervals taken together does not sum to unity, as the probability calculus demands, but it is infinite.

Tempting as this improper distribution may be, it suffers the same problem as the proper distribution. It still adds physical properties to the dome. For a consequence of it is that spontaneous motion in time $t=1$ to $t=2$ has probability ϵ and spontaneous motion in the time interval $t=2$ to $t=4$ is 2ϵ . Motion in the one interval is twice as probable--no more no less--than motion in the other. But nothing in the physics licenses this precise judgment. All the physics says is that motion in each interval is "possible"—and that is all. Once again, we have passed from being inductive logicians to being physicists, adding more physical properties to the system than Newton's theory has already given it.¹⁵

¹⁵ Some cannot resist the temptation of representing our belief as a set of probability measures, for example, the set of measures associated with all $0 < \tau \leq \infty$. This proposal succeeds only in so far as it recapitulates the austere logic described in Section 6 below. My principal objection is that it amounts to using an additive measure merely to simulate a belief distribution that is inherently non-additive and thus inherently non-probabilistic. That such simulation is possible does not show that our beliefs are really probabilities after all. No doubt, with enough ingenuity, we could use any other sort of measure to simulate the austere logic of Section 6.

6. How to Infer Inductively about Indeterministic Systems

Something has gone very wrong and, in my view, what has gone wrong is quite simple. We are trying to force the wrong inductive logic onto the dome. How can we select the right one? The material theory of induction directs us to look to the prevailing material facts. They will fix the inductive logic, just as the probabilistic law of radioactive decay led us to a probabilistic logic for the radioactive decay of an atom. The physics of the dome is more impoverished than that of radioactive decay. So we should expect a more impoverished logic.

Once this notion is accepted, it becomes a somewhat mechanical exercise to read the relevant inductive logic from the physics. For radioactive decay, the chance of decay in time $t=5\tau$ is 0.99; so our degree of belief in that decay is 0.99. However the indeterministic physics of the dome does not give us real-valued degrees. It actually says rather little. It just says that a spontaneous motion in this or that time is possible; and that it all. It provides no degrees of possibility: not 50% possible, not 95% possible; and no comparative measures: not more possible, less possible, twice as possible. It just asserts what is possible; and by logical implication we can also know what is necessary and impossible. These three assignments, necessary, possible and impossible, become the three values of our inductive logic. The translation of the material facts in the physics to the inductive logic is illustrated in the table:

<i>What the physics says:</i>	<i>What it induces in the inductive logic:</i>
The present state does not fix the future (indeterminism). The physics just tells us that a future state is <i>necessary, possible or impossible</i> .	The inductive logic for the support [A B] of A from B has three values: <i>nec, poss, imp</i> .
If the motion happens in time [10,20), then it <i>necessarily</i> happens in [0,100).	[H(0,100) H(10,20)] = <i>nec</i>
Motion in any later non-zero interval is <i>possible</i> , given E: the mass is at rest at the apex of the dome at t=0.	[H(0,10) E] = [H(0,100) E] = [H(10,20) E] = ... = <i>poss</i>
If the motion happened in [0,10), it is <i>impossible</i> in [20,30).	[H(20,30) H(0,10)] = <i>imp</i>

Table 1. Material Facts Dictate an Inductive Logic

The table has given a few obvious illustrations of a more general system. It is easy to see that the full logic is generated by simple rules. They are:

The complete inductive logic of the dome

[A|B] = *nec*, if B entails A

= *imp*, if B entails not A

= *poss*, otherwise

7. Complete Ignorance

The natural inductive logic of the dome is very weak. Indeed one may want to say that we are in complete ignorance over the time of the spontaneous motion. While we may say that informally, there is a precise sense in which it is true of the logic just recovered. Elsewhere (Norton, manuscript b) I have urged that the literature in probability and inductive inference has long harbored serviceable instruments for determining the properties of a belief distribution that coincides with complete ignorance. These instruments are platitudes of evidence. Yet we have misdiagnosed them as paradoxical since they require belief distributions that are not probability distributions.

7.1 Principle of Indifference

This first instrument was made famous by its association with the classical interpretation of probability. It asserts:

(PI) *Principle of Indifference*. If we are indifferent among several outcomes, that is, if we have no grounds for preferring one over any other, then we assign equal belief to each.

The familiar paradoxes of indifference stem from the fact that, in cases of far-reaching ignorance, we will be indifferent over many distinct partitions of the outcome space. For example, if two balls can each be red “R” or black “B,” we may be indifferent over the outcomes

two-R, one-R, no-R;

or we may be indifferent over the outcomes

RR (=first ball R, second ball R), RB, BR, BB.

In the first, we assign equal belief to each of two-R and one-R; in the second we assign equal belief to each of RR, RB and BR. But since RR is just the same as two-R, our analysis requires us to assign equal belief to one-R and each of its two disjunctive parts RB and BR. That is impossible if the beliefs are (non-zero) probabilities.

Other examples are similar but more elaborate.¹⁶ They all end up establishing that a thorough application of the principle in cases of complete ignorance forces us to assign the same belief to a contingent outcome and to each of its contingent disjunctive parts. It is a near universally held judgment in the literature in probability theory that this is what the Principle of Indifference demands. This demand is paradoxical if we presume that the epistemic state of

¹⁶ For example, we may be indifferent over where a real number x lies in 0-100. So we assign equal belief to x lying in 0-50 or 50-100. If, however, we assign indifference over the x^2 , then we assign equal weight to x^2 lying in each of 0-2500, 2500-5000, 5000-7500, 7500-10000.

Proceeding as before, we end up assigning equal belief to x^2 lying in the interval 2500-10000 and each of its disjunctive parts, 2500-5000, 5000-7500, 7500-10000.

complete ignorance can be represented by a probability measure, for no probability measure can have that property.¹⁷

The standard resolution has been to disparage the principle as defective, even though it is a platitude of evidence: that beliefs should not differ if there are no grounds for the difference. The better resolution is discard the presumption that the epistemic state of complete ignorance can be represented by a probability measure. Then the demands of the principle are readily met. The logic of the dome, for example, does allow us to assign equal belief to a contingent proposition and each of its disjunctive parts. Take the interval $[T_1, T_3) = [T_1, T_2) \cup [T_2, T_3)$, where $T_1 < T_2 < T_3$. In the logic of the dome, we have equal degrees of support on E assigned to each

$$[H(T_1, T_3)|E] = [H(T_1, T_2)|E] = [H(T_2, T_3)|E] = \textit{poss}$$

Thus the belief distributions of the dome are the ignorance distributions naturally picked out by the Principle of Indifference.

7.2 Principle of Invariance of Ignorance

The second principle is that an ignorance distribution should remain unchanged if we redescribe the outcomes in a symmetrical way. The simplest application is a derivation of the principle of indifference. If we are really indifferent to outcomes A, B and C, then our belief will be the same over outcomes A', B' and C', where these are any permutation of A, B and C. For example: A'=B, B'=A, C'=C; or A'=B, B'=C, C'=A; etc. One readily sees that uniformity of belief is the only belief distribution that is unchanged under these permutations. The principle, in its strictest and most defensible form, (Norton, manuscript, §2.1) is

(PII) *Principle of the Invariance of Ignorance.* An epistemic state of ignorance is invariant under a transformation that relates symmetric descriptions.¹⁸

¹⁷ The obvious escape is to seek the finest partition and apply the principle there only. If the space is spanned by a continuous variable, there is no finest partition. Even if the space is finite—outcomes A, B or C—we cannot preclude refinement of the space by introducing a new outcome “X” from another other space. So the refinement is $A \vee X, A \vee \neg X, B, C$.

¹⁸ Symmetric descriptions are defined here as pairs of descriptions meeting two conditions:

This requirement of invariance has been used extensively in the objective probability literature. What is less quickly acknowledged is that it leads to paradoxes in probability theory just as readily as the Principle of Indifference. The mechanism that generates the paradoxes is analogous. Each invariance constrains the belief distribution. Many invariances, such as arise when one is truly ignorant, can so overconstrain the belief distribution that it cannot be a probability distribution.

As a simple example, imagine that we are completely ignorant as to the value of some real number x in the interval $(0,1)$. That means that our belief distribution must be invariant under the self-inverting transformation $x'=1-x$. Many more such self-inverting transformations are possible, such as¹⁹ $x' = 1 - (2x-x^2)^{1/2}$. It is easy to show (Norton, manuscript, §3.2) that invariance under all of them forces a probability distribution to be everywhere zero.

As before the Principle of Invariance of Ignorance is a platitude of evidence that merely requires that we not harbor differences of belief among two descriptions if we have no grounds for picking them apart. If probability measures cannot accommodate the platitude, then we must infer that ignorance distributions of belief cannot be probability measures.

Once we forgo the assumption that ignorance distributions must be probability distributions, the Principle of Invariance of Ignorance will pick out a unique ignorance distribution of belief. The key intuition is the following. Assume that we are in complete ignorance concerning contingent outcomes in a space described by the (not necessarily mutually exclusive) contingent propositions A_1, \dots, A_n . That distribution would be unaffected if we had inadvertently mislabeled the propositions by their negations, so whenever we contemplated A_j , we would actually be contemplating $\neg A_j$. Imagine that we are in complete ignorance of the truth

(S1) The two describe exactly the same physical possibilities; and each description can be generated from the other by a relabeling of terms, such as the additional or removal of primes, or the switching of words.

(S2) The transformation that relates the two descriptions is “self-inverting.” That is, the same transformation takes us from the first description to the second, as from the second to the first.

¹⁹ $x'=f(x)=1-x$ is self inverting since $f(f(x))=x$. Similarly $x'=g(x)=1-(2x-x^2)^{1/2}$ is self inverting in that $g(g(x))=x$.

of Julius Caesar's G: "Gallia est omnis divisa in partes tres," which is translated as "All Gaul is not divided into three parts." Our epistemic state regarding G would not change if I now admit that the translation just given is incorrect and should be "All Gaul is divided into three parts."

More formally, the negation map that takes A_i to $A'_i = -A_i$ is a self-inverting transformation that must leave the distribution of belief unchanged. It is an easy but somewhat tedious exercise (Norton, manuscript b, §6.3) to demonstrate that the unique monotonic belief distribution satisfying this invariance is one that assigns the same degree of belief to all contingent propositions.²⁰ That distribution once again coincides with the belief distribution arising in the logic of the dome, in the sense that $[H(T_1, T_2)|E]$ has the same value for all admissible T_1 and T_2 .

8. Conclusion

The inductive logic of indeterministic systems such as the dome is non-probabilistic. For that is the sort of logic that the material facts of the system require. That the material facts prescribe the inductive logic is, in my view, not just true for these special systems, but for all systems over which we infer inductively. One may doubt that the claim can be so easily generalized to all systems. Let me address and answer two hesitations that may motivate these doubts.

First, one may doubt that the odd logic just revealed has anything to do with our world, which, at least on the level of everyday experience, seems free of the indeterminism exemplified

²⁰ Monotonicity is the demand that the degree of support accorded some outcome should be no less than that accorded its logical consequences. The invariance is required only for the degree of belief assigned to contingent propositions. For, under the negation map, necessary truths, to which full belief should be assigned, are mapped to necessary falsehoods, to which full disbelief is assigned. One could extend the analysis to include these non-contingent propositions by dropping the assumption that we are able to identify the logical truths of the system and are as ignorant of them as of the contingent propositions. However that is a less interesting case.

by the dome.²¹ As any parent of a three-year old knows, things do not spring into motion in ordinary experience unless they are pushed. While that may be true, it does tacitly concede my major point: that the inductive logic applicable in some domain does depend on the facts that prevail there.²² That inductive logic in our world of common experience may not be the same as that of the dome is merely a reflection of the differing facts that prevail in the two domains. That is just what the material theory of induction claims. Once we accept that our logic depends upon at least some facts, are we so sure that the only pertinent facts are exotic ones like the absence of indeterministic systems? Might not other less exciting facts also be pertinent to our selection of the right inductive logic so that multiple inductive logics might well be appropriate to different domains in our ordinary experience?

Second, one may worry that the analysis is only applicable to systems whose physical properties are fully specified. Why should the analysis also apply to more realistic, real-world inductive inference problems in which the full properties of the system are unknown? I do admit that the example is a little contrived. That was the price paid for an example in which we have complete control of all relevant facts that govern the ways that uncertainty enters into our analysis. For a complete control of those facts makes it quite easy to see just which inductive logic is appropriate. In real world cases of inductive inference, the facts that govern the uncertainties are less clear to us. However that sort of muddiness is no reason to think that things are any different. There still are facts governing our uncertainties and they will still dictate the appropriate logic. However now it will be harder for us to see precisely what that logic is because we are uncertain of the pertinent facts. In real world cases, we do struggle to see just how much this piece of evidence does lend support to that hypothesis. That is just what we would expect from the material theory of induction when we are unsure of the facts that govern the uncertainties. If one does not subscribe to the material theory of induction and holds out the hope for a universal logic of induction applicable everywhere, this continuing elusiveness will be

²¹ There may be non-probabilistic indeterminism in our world. Standard big bang cosmology allows many different matter densities to emerge from the big bang singularity without assigning probabilities to the different possibilities.

²² Contrast this with deductive logic, which is universal. The pathology of the dome does not preclude our use of deductive logic in reasoning about it.

a chronic frustration. That in turn, it seems to me, goes a long way to explain why induction has perennially seemed such a murky topic.

References

- Alper, Joseph S., Bridger, Mark, Earman, John and Norton, John D. (2000) "What is a Newtonian System? The Failure of Energy Conservation and Determinism in Supertasks," *Synthese*, 124, pp. 281-293.
- de Finetti, Bruno (1937) "Foresight: Its Logical Laws, Its Subjective Sources," trans. H. E. Kyburg, in H. E. Kyburg and H. Smokler, eds., *Studies in Subjective Probability*. New York: John Wiley & Sons, 1964, pp. 93-158.
- Norton, John D. (1999) "A Quantum Mechanical Supertask" *Foundations of Physics*, 29, pp. 1265-1302.
- Earman, John (1992) *Bayes or Bust*. Cambridge, MA: Bradford-MIT.
- Jaynes, E. T. (2003) *Probability Theory: The Logic of Science*. Cambridge: Cambridge University Press.
- Lewis, David (1980) "A Subjectivist's Guide to Objective Chance," pp. 263-93 in R.C. Jeffrey ed., *Studies in Inductive Logic and Probability*. Berkeley: University of California Press.
- Norton, John D. (2003) "A Material Theory of Induction" *Philosophy of Science*, 70, pp. 647-70.
- Norton, John D. (2003a) "Causation as Folk Science," *Philosophers' Imprint*, 3(No.4); www.philosophersimprint.org/003004/
- Norton, John D. (2005) "A Little Survey of Induction," in P. Achinstein, ed., *Scientific Evidence: Philosophical Theories and Applications*. Johns Hopkins University Press, 1905. pp. 9-34.
- Norton, John D. (forthcoming), "Probability Disassembled" *British Journal for the Philosophy of Science*.
- Norton, John D. (manuscript a), "The Dome: An Unexpectedly Simple Failure of Determinism." www.pitt.edu/~jdnorton
- Norton, John D. (manuscript b), "Ignorance and Indifference." www.pitt.edu/~jdnorton
- Savage, Leonard J. (1972) *The Foundations of Statistics*. Second revised ed. New York: Dover.
- Smith, C. A. B. (1961) "Consistency in Statistical Inference and Decision," *Journal of the Royal Statistical Society, Series B*, 23, pp. 1-25.