Induction and Probability

Alan Hájek and Ned Hall

In *The Blackwell Guide to the Philosophy of Science,* eds. Peter Machamer and Michael Silberstein, Blackwell, 149-172.

We will discuss induction and probability in that order, aiming to bring out the deep interconnections between the two topics; we will close with a brief overview of cutting-edge research that combines them.

1. Induction: some preliminaries

Arguably, Hume's greatest single contribution to contemporary philosophy of science has been the problem of induction (1739). Before attempting its statement, we need to spend a few words identifying the subject matter of this corner of epistemology. At a first pass, induction concerns *ampliative* inferences drawn on the basis of evidence (presumably, evidence acquired more or less directly from experience)—that is, inferences whose conclusions are not (validly) entailed by the premises. Philosophers have historically drawn further distinctions, often appropriating the term "induction" to mark them; since we will not be concerned with the philosophical issues for which these distinctions are relevant, we will use the word "inductive" in a catch-all sense synonymous with "ampliative". But we will follow the usual practice of choosing, as our paradigm example of inductive inferences, inferences about the future based on evidence drawn from the past and present.

A further refinement is more important. Opinion typically comes in *degrees*, and this fact makes a great deal of difference to how we understand inductive inferences. For while it is often harmless to talk about the conclusions that can be rationally *believed* on the basis of some evidence, the force of any given evidence is more accurately measured by noting its effects on the rational assignment of *degrees* of belief. The usual assumption—one that directly connects the two topics of this chapter—is that rationally warranted degree of belief can be modeled as a species of *probability*, inductive inference itself being modeled by some rule for changing probabilities in light of evidence.

The strength of the support that some evidence gives some hypothesis is *not* measured by the degree of belief in the hypothesis that is warranted in light of that evidence. There is widespread disagreement among philosophers of science as to how exactly it should be measured (see Milne 1996); but all parties agree that the notion is an essentially *comparative* one, in that it depends not only on how likely an hypothesis is, in light of the given evidence, but also on how likely it would be, in light of *different* evidence.

2. The first problem of induction

One serious problem that inductive inferences present us with is that of saying with any precision what distinguishes the *good* ones from the *bad* ones. We will get to that problem shortly, and because we take it up second will call it the "second problem of induction". It should be distinguished from the first problem, which is that of saying why the 'good' ones *deserve* this label.

For various reasons (whose elucidation space does not permit), we think it is best to formulate this problem as a problem about settling a conflict between two rival inductive methods—rival sets of rules for adjusting degrees of belief in light of evidence. Let us personify two such rivals in the form of Billy and Suzy, two friends. Suzy is a paragon of cognitive virtue: she always evaluates the impact of evidence on hypotheses in accordance with the "right" inductive principles. Billy evaluates the force of evidence in accordance with different principles, a fact which shows up in the following bizarre behavior: He regularly sticks his fingers in light sockets, always getting a nasty shock when he does so. To be clear, Billy doesn't *like* these shocks *at all*. It's just that each time, his inductive methods lead him to the conclusion, given the evidence available to him, that it is overwhelmingly likely that the sensation will be exquisitely pleasant. Billy and Suzy disagree about these predictions, and since—let us stipulate—they have exactly the same evidence, their disagreement traces to the different principles they adhere to in evaluating the force of such evidence. Is it possible to provide any compelling argument for the conclusion that the inductive principles to which Suzy adheres are rationally superior to Billy's?

There are two relevant parties here, and we need to consider the possibility that there is an argument compelling to one but not the other. Let us stipulate that such an argument must make use of acceptable premises that do not beg the question against the party to be compelled. We will take it that such premises will at least include all propositions detailing the evidence available to Billy and Suzy. Let us further stipulate that the premises must, in some sense, *support* the given conclusion, and that they can do so in one of only two ways: either they entail the conclusion that Suzy's inductive principles are rationally preferable to Billy's, or they provide some measure of inductive support for this conclusion.

It might seem that no argument of the first kind that would be compelling to *either* Billy or Suzy is possible, especially if we limit our attention to arguments that proceed only from the available evidence, and that attempt to establish the superiority of Suzy's inductive methods over Billy's by way of the intermediate conclusion that her methods will in the future yield correct predictions more often than Billy's. Here, the point, familiar since Hume, that the past places no logical constraints on the future renders such an intermediate conclusion inaccessible. But there is always the possibility of finding additional, non-question-begging premises, or of finding some other route to the conclusion—loopholes that, as we'll see below, Reichenbach's "pragmatic" justification of induction attempts to exploit.

Moreover, there is at least one clear way in which such an argument could be constructed: namely, if Billy's inductive rules undermine themselves by predicting, given the evidence, that they will systematically issue *false* predictions in the future. If Suzy's principles do not undermine themselves in this way, then they will clearly be rationally preferable; what's more, this conclusion validly follows from premises perfectly acceptable to both. Still, there is little point in hoping for such an argument, as it turns out to be far too easy—and costless—to construct inductive methods that are immune from self-undermining.[1] So we might as well build into our description of the Billy/Suzy scenario that Billy is adhering to just such a method.

3. The inductive justification of induction

What then of the possibility of a compelling *inductive* argument? None could succeed in convincing Billy. For either the evidence available to Billy inductively warrants, by his lights, the conclusion that Suzy's inductive methods are preferable to his own, or it doesn't. But if the evidence warrants this conclusion then his principles undermine themselves, and we have stipulated that they are immune from such self-undermining. So consider whether some inductive argument could be produced that would provide *Suzy* with compelling grounds for holding her inductive methods to be rationally preferable to Billy's.

To give her the best possible case, let us suppose that Suzy's own inductive principles strongly endorse, in light of the evidence available to her, the conclusion that those very

---

[1] Space prevents us from giving the details, but the key idea is to build into one's inductive principles a version of van Fraassen's Reflection Principle, which we will discuss below.

principles will yield wildly successful predictions in the future, whereas Billy's will yield an unbroken string of falsehoods. It seems, then, that she has a compelling and powerful argument for the target conclusion.

But there has been a bit of sleight of hand. The problem is *not* that her "inductive justification of induction" is circular or question-begging, for given that *she* is its target it manifestly *isn't*. (See van Cleve 1984; any lingering sense that it *is* can be explained away by noting that no such inductive argument could convince *Billy*.) Rather, we need to remember that strength of inductive support is a comparative notion. In the case at hand, the "track record" of Suzy's inductive methods provides, by the lights of these methods, *extra* reason to have faith in them only to the extent that other track records were possible that would have yielded a more pessimistic prediction—i.e., only if it was at least possible for the evidence to produce a self-undermining verdict. But it is prima facie quite desirable to adhere to an inductive method that is immune from the possibility of self-undermining, particularly given that this is both easy and costless to do. Assuming that Suzy *is* following such a method, the track record of its successes, however spectacular, contributes nothing *at all* to the case in its favor.

In any case, even at the beginning of inquiry—before any evidence has been amassed—there remains just as clear and intuitive a difference in the rational acceptability of Suzy's and Billy's inductive methods, one which need wait upon no "supporting evidence" in order to become visible: The difference is that hers are superior. Not, then, because the evidence favors them.

4. The pragmatic justification

It seems transparent that no valid argument could be produced whose premises are obviously correct and non-question-begging and whose conclusion is that our particular inductive practices are rationally warranted. Famously, Reichenbach (1938 and 1949) attempts to produce just such an argument.

The more formal version of the argument seeks to justify the use of the "straight rule" in making predictions about the limiting relative frequency of outcomes in an infinitely repeated experiment (the rule predicts this limit to equal the frequency observed so far); it has been so thoroughly discussed in the literature that we will pass over it (see Salmon 1967). The less formal version is more clever, and less talked-about. For the sake of definiteness, let us use the label "the scientific method" (SM) as a name for whatever method it is we use—at least, when we are at our cognitive best—for drawing inductive conclusions. And let us suppose that we have some standard of success for an inductive method—say, that the long-range frequency of correct inferences drawn on its basis must be sufficiently high. For any given inductive method, there is of course no guarantee that it will succeed in this sense: the world must cooperate, and it is a contingent matter whether it will do so. Reichenbach thus grants that we can have no a priori guarantee of the success of SM. But he argues that we are "pragmatically" justified in our continued adherence to SM, since if *any* method will succeed, *it* will. For suppose the world is such that some rival method RM will succeed, but SM will not. Well, a central component of SM consists in projecting observed regularities into the future, and in a world where RM succeeds, a relevant such regularity simply *is* the pattern of RM's successes. If, for example, a successful method for predicting the future in world w is to consult an oracle, then SM will eventually establish that the oracle is reliable—and so SM will itself ultimately yield the advice that one should consult the oracle when making predictions. The key claim—that if the world is nice enough to allow for the success of some inductive method, then it is nice enough to allow for the success of SM—is simply a generalization of this example. So we have a demonstrative argument that SM will be successful if any method will; hence, it seems, a demonstrative argument that we are rationally warranted in adhering to SM.

But the example of the oracle is entirely misleading. For a long string of successful predictions by such an oracle surely constitutes a very *salient* regularity. Suppose the success of the rival inductive method is not nearly so visible; why should we have any confidence that SM will "latch onto" this string of accurate predictions?

We should have *no* such confidence if the following condition holds: For every proposition A about the future, there are rival inductive methods that have been highly successful, and equally successful, in the past, but that disagree widely as to the likely truth of A (given the available evidence). If that condition is met, then the argument for the crucial premise fails disastrously—for *which* string of successes should SM latch onto? It is easy to see what went wrong: the argument involved a bit of misdirection in getting is to agree implicitly that SM was, in the imagined world w, up against just *one* rival RM. More plausible is that it would be up against a battery of rivals so extensive that they fail as a group to agree on any prediction of substance. It is wishful thinking to suppose that SM could somehow pick the "winner".

5. Hume's 'skeptical solution'

Hume's own "skeptical" solution to his problem of induction foreshadowed an important movement in contemporary epistemology, which seeks to "naturalize" the subject (Kornblith 1985). For Hume took it that no rational basis for induction is possible, while adding that perfectly legitimate empirical psychological questions remain about how exactly it is that deliberating agents draw inductive conclusions from evidence. Hume's own answer emphasized the central role of the apparently brute psychological disposition he called "custom" or "habit"; contemporary fans of this kind of naturalized approach to inductive epistemology could presumably be expected to draw on much more sophisticated theories of human cognitive psychology.

A serious worry is that it is unclear that the naturalizing move in epistemology leaves room for a legitimate, coherent sub-discipline of *normative* epistemology, a discipline that seeks to articulate the principles according to which we *ought* to form our beliefs. That is unfortunate, since our natural or "untutored" cognitive abilities in the inductive domain are notoriously and systematically unreliable, particularly when we're in a situation that forces us to attend sensitively to the probabilistic bearing of evidence on hypotheses (see Kahneman, Slovic, and Tversky 1982). It would seem to require careful *a priori* reflection to distinguish rational inductive inferences from mistakes. At the very least, the defender of a purely naturalized epistemology of induction owes us an account of how *else* we might systematically identify and guard against inductive error.

6. Popper's falsificationism

Popper (1968) argues for a different way of dismissing the problem of induction: while agreeing with Hume that no rational justification of induction can be found, he insists that this result is innocuous, simply because induction forms no part of the practice of science. According to Popper, scientists propose "conjectures", and then subject these conjectures to severe observational tests in an effort to falsify them. He claims that we are never rationally warranted in considering such an hypothesis to be probable, given that it has passed such tests. And since, according to the simplest version of falsificationism, deductive relations are all that we need attend to in order to check that an hypothesis has been refuted by some evidence, the problem of induction poses no threat to the rationality of scientific practice.

As a descriptive claim about what scientists, *qua* scientists, actually *do*—let alone about what they *believe* about what they do—Popper's view strikes us as absurd. But even as a normative claim it fares little better. The simplest and most devastating point was nicely emphasized by Putnam (1974): Popper seems willfully blind to the fact that we use evidence from the past and present as a basis for making practical *decisions*, decisions whose rationality is hostage to the rationality of the inferences drawn about their likely consequences on the basis of the given evidence. What would Popper say, for example, about the disagreement between Billy and Suzy? That Billy's behavior is somehow rationally permissible, even in light of the extensive and painful evidential record? Even were it understood merely as a claim about the rationality of belief, falsificationism would be unpalatable; but belief and action are too inextricably linked to sustain such an understanding. The consequences of the position for the rationality of decision provide it, ironically enough, with a decisive refutation.

7. The dogmatic response

The final approach to the first problem of induction that we will consider we call the "dogmatic" response—not as an insult, but because the word nicely summarizes its main features. For according to the dogmatic response, induction is perfectly rational—certain ways of adjusting degrees of belief in the light of evidence are rationally warranted, and certain other ways are irrational—but absolutely no justification can be given of this claim, not even a justification of the kind that would only be compelling to the likes of Suzy (see Strawson 1952). It is rather that the fact that certain inductive inferences are rational and certain others irrational (and perhaps still others neither rational nor irrational) is a brute epistemological fact, incapable of further philosophical explanation or defense.

The principal merits of the view are clear enough. It allows us to maintain, contra Hume and other skeptics about induction, a vigorous distinction between rational and irrational inductive methods and inferences, and it acquires at least some measure of plausibility from the dismal failure of more ambitious attempts to give a justification of induction. Still, the view should only be seen as a kind of philosophical last resort. For there are too many interesting questions about which the dogmatic response falls silent. Notably, whereas any attempt to provide a substantive justification of induction can be expected, to the extent that it succeeds, also to provide insights into what distinguishes good from bad inductive inferences, the dogmatic response is hopeless in this regard. And, as noted above, the problem of providing a clear explication of the distinction between rational ('good') and irrational ('bad') inductive inferences is a deep and central one in its own right. We turn now to a brief discussion of some of the main philosophical approaches in this area.

8. The second problem of induction: Syntactic approaches

Traditionally, logic aims to distinguish valid from invalid arguments by virtue of the syntactic form of the premises and conclusion (e.g., any argument that has the form *p and q, therefore p* is valid in virtue of this form). But the distinction between valid and invalid is not fine enough: after all, many invalid arguments are perfectly good, in the sense that the premises provide strong inductive support for the conclusion. Carnap (1950) described this relation of support as the *logical probability* that an argument's conclusion is true, given that its premises are true—hoping that logic, more broadly conceived, could give it a *syntactic* analysis. We will discuss Carnap's approach in more detail below. Other, less ambitious approaches tried to find syntactic criteria for "qualitative confirmation"—criteria, that is, that would identify at least some instances in which evidence raised the probability of an hypothesis, to at least some degree. (See Hempel 1945a,b for an excellent overview of work in this area.)

These attempts to design a "logic" of induction on the model of formal deductive logic did not succeed. The decisive problem concerns the language-dependence that any such "logic" would have to exhibit. Consider, for example, a language used to represent the outcomes of random draws from an urn filled with colored balls; let the language contain the color predicates "blue" and "green", and also, in the spirit of Goodman (1983), the predicate "grue", where "x is grue at draw i" is equivalent to "x is green at draw i and $i \leq 1{,}000{,}000$ or x is blue at draw i and $i > 1{,}000{,}000$". Surely the "logical strength" of the argument, "the first million draws are green, therefore the next draw will be green" is greater than 1/2; if syntax is all that matters, then so too is the logical strength of the argument, "the first million draws are grue, therefore the next draw will be grue". But the two conclusions contradict each other, and so cannot both receive probability greater than 1/2.

One might try to specify a canonical language, to sentences of which the syntactic rules, whatever they are, are meant to apply—a language free of such monstrosities as "grue". But not only does traditional logic find no need for such a procedure, it is also extraordinarily difficult to see how one could carry it out, at least if we want to analyze the inductive strength of any argument of real interest. By the middle of the 17th century, the available evidence strongly supported Keplerian over Ptolemaic astronomy; but what would be the canonical language in which to translate this evidence and these hypotheses, in order to analyze the differential support *syntactically*?

9. The second problem: Modest probabilism

One might agree with Carnap that induction should be modeled using the tools of probability theory, while denying that syntactic analysis alone can provide or even constrain the values of the relevant probability function. And indeed, what we will call "modest probabilism" about induction and confirmation has become increasingly popular since the demise of logical empiricism. We call this approach "probabilism" because it sees the inductive support or degree of confirmation that evidence E gives hypothesis H as measured by somehow comparing P(H), the probability of H, with P(H|E), the conditional probability of H, given E. (We will have more to say about these quantities shortly.) Perhaps this inductive support is measured by the difference P(H|E) – P(H); perhaps by the ratio P(H|E)/P(H); perhaps in some other way. (See Good 1985.) But the approach is modest to the extent that it is agnostic about the nature or source of the "confirmation-probability" in question. Its agnosticism notwithstanding, modest probabilism is able to achieve some remarkable successes. For example, it explains straightaway the success (such as it is) of the hypothetico-deductive account of confirmation. For if H implies E, and if P(E) < 1, then it follows at once that P(H|E) > P(H) (for this condition is equivalent to P(E|H) > P(E)). More interestingly, modest probabilism neatly explains away the Raven's Paradox, and can be easily adapted to illuminate the confirmation of hypotheses that are themselves probabilistic (See Earman 1992 for a fuller discussion.)

Partly in order to flesh out the resources of and problems for this probabilistic approach, we will now switch gears slightly, and take up the second of our topics: an investigation of probability theory and the most important attempts at explicating its conceptual foundations. We begin with an overview of the widely accepted *mathematical* foundations.

10. Kolmogorov's axiomatization

Probability theory was inspired by games of chance in 17th century France and inaugurated by the Fermat-Pascal correspondence. However, its axiomatization had to wait until Kolmogorov's classic book (1933). Let Ω be a non-empty set ('the universal set'). A *sigma-field* (or *sigma-algebra*) on Ω is a set $\mathcal{F}$ of subsets of Ω that has Ω as a member, and that is closed under complementation (with respect to Ω) and countable union. Let P be a function from $\mathcal{F}$ to the real numbers obeying:

1. $P(A) \geq 0$ for all $A \in \mathcal{F}$.

2. $P(\Omega) = 1$.

3. $P(A \cup B) = P(A) + P(B)$ for all $A, B \in \mathcal{F}$ such that $A \cap B = \varnothing$.

Call P a *probability function*, and $(\Omega, \mathcal{F}, P)$ a *probability space*.

We could instead attach probabilities to members of a collection of *sentences* of a formal language, closed under truth-functional combinations.

It is controversial whether probability theory should include Kolmogorov's further axiom:

4. (Continuity) $E_n \downarrow \varnothing$ implies $P(E_n) \to 0$ (where $E_n \in \mathcal{F}\ \forall n$)

Equivalently, we can replace the conjunction of axioms 3 and 4 with a single axiom:

3′. (Countable additivity) If $\{A_i\}$ is a countable collection of (pairwise) disjoint sets, each $\in \mathcal{F}$, then

$$P\left(\bigcup_{n=1}^{\infty} A_n\right) = \sum_{n=1}^{\infty} P(A_n)$$

The conditional probability of X given Y is standardly given by the ratio of unconditional probabilities:

$$P(X|Y) = \frac{P(X \leftrightarrow Y)}{P(Y)} \text{, provided } P(Y) > 0.$$

We can now prove versions of *Bayes' theorem*:

$$P(A|B) = \frac{P(B|A).P(A)}{P(B)}$$

$$= \frac{P(B|A).P(A)}{P(B|A).P(A) + P(B|\neg A).P(\neg A)}$$

More generally, suppose we have a partition of hypotheses $\{H_1, H_2, ..., H_n\}$, and evidence E. Then we have, for each i:

$$P(H_i|E) = \frac{P(E|H_i)P(H_i)}{\sum_{j=1}^{n} P(E|H_j)P(H_j)}.$$

The $P(E|H_i)$ terms are called *likelihoods*, and the $P(H_i)$ terms are called *priors*.

If $P(X|Y) = P(X)$—equivalently, if $P(Y|X) = P(Y)$; equivalently, if $P(X \cap Y) = P(X)P(Y)$—then X and Y are said to be *independent*. Two cautions: firstly, the locution 'X is independent of Y' is somewhat careless, encouraging one to forget that independence is a relation that events or sentences bear *to a probability function*. Secondly, this technical sense of 'independence' should not be identified unreflectively with causal independence, or any other pretheoretical sense of the word, even though such identifications are often made in practice. If $P(X|Y) > P(X)$—equivalently, if $P(Y|X) > P(Y)$—then X and Y are *positively correlated*. A cornerstone of any probabilistic approach to induction is the idea that evidence about the observed is positively correlated with various hypotheses about the unobserved.

We now turn to the so-called '*interpretations*' of probability. The term is misleading twice over. Various quantities that intuitively have nothing to do with 'probability' obey Kolmogorov's axioms—for example, length, volume, and mass, each suitably normalized—and are thus 'interpretations' of it, but not in the intended sense. Conversely, the majority of the most influential 'interpretations' of P violate countable additivity, and thus are not genuine interpretations of Kolmogorov's full probability calculus at all. Be that as it may, we will drop the scare quotes of discomfort from now on.

12. The classical interpretation

The classical interpretation, which owes its name to its early and august pedigree, (notably the *Port-Royal Logic*, Arnauld 1662, and Laplace 1814) purports to determine probability assignments in the face of no evidence at all, or symmetrically balanced evidence. In such circumstances, probability is shared equally among all the possible outcomes, so that the classical probability of an event is simply the fraction of the total number of possibilities in which the event occurs—a version of the so-called *principle of indifference*. Unless more is said, it is also arguably the interpretation furthest removed from considerations of induction, reflecting as it does a certain *a prioristic* innocence: in typical applications, the number of possibilities, and thus the share that each gets of the total probability, remain the same (e.g. 3/6) whatever the outcomes in the actual world happen to be. Unfortunately, the classical interpretation can apparently yield contradictory results when there is no single privileged set of possibilities, as Bertrand (1989) brought out in his paradoxes. Classical probabilities are only finitely additive (see de Finetti 1974).

13. The logical interpretation

Logical theories of probability retain the classical interpretation's guiding idea that probabilities can be determined a priori by an examination of the space of possibilities. However, they generalize it in two important ways: the possibilities may be assigned *unequal* weights, and probabilities can be computed whatever the evidence may be, symmetrically balanced or not. Indeed, the logical interpretation, in its various guises, seeks to codify in full generality the degree of support or confirmation that a piece of evidence E confers upon a given hypothesis H, which we may write as $c(H, E)$.

Early proponents of logical probability include Keynes (1921), W.E. Johnson (1932), and Jeffreys (1939). However, by far the most systematic study of logical probability was by Carnap. His formulation of logical probability begins with the construction of a formal language. In (1950) he considers a class of very simple languages consisting of a finite number of logically independent monadic predicates (naming properties) applied to countably many individual constants (naming individuals) or variables, and the usual logical connectives. The strongest (consistent) statements that can be made in a given language describe all of the individuals in as much detail as the expressive power of the language allows. They are conjunctions of complete descriptions of each individual, each description itself a conjunction containing exactly one occurrence (negated or unnegated) of each predicate of the language. Call these strongest statements *state descriptions.*

Any probability measure $m(–)$ over the state descriptions automatically extends to a measure over all sentences, since each sentence is equivalent to a disjunction of state descriptions; m in turn induces a confirmation function $c(–,–)$:

$$c(H, E) = \frac{m(H \,\&\, E)}{m(E)}$$

There are obviously infinitely many candidates for *m*, and hence *c*, even for very simple languages. Carnap argues for his favored measure "*m**" by insisting that the only thing that significantly distinguishes individuals from one another is some qualitative difference, not just a difference in labeling. A *structure description* is a maximal set of state descriptions, each of which can be obtained from another by some permutation of the individual names. *m** assigns each structure description equal measure, which in turn is divided equally among their constituent state descriptions. It gives greater weight to homogenous state descriptions than to heterogeneous ones, thus 'rewarding' uniformity among the individuals in accordance with putatively reasonable inductive practice. It can be shown that the induced *c** allows inductive learning from experience—as, annoyingly, do infinitely many other candidate confirmation functions. Carnap claims that *c** nevertheless stands out for being simple and natural.

He later generalizes his confirmation function to a continuum of functions c_λ. Define a *family* of predicates to be a set of predicates such that, for each individual, exactly one member of the set applies, and consider first-order languages containing a finite number of families. Carnap (1963) focuses on the special case of a language containing only one-place predicates. He lays down a host of axioms concerning the confirmation function *c,* including those induced by the probability calculus itself, various axioms of symmetry (for example, that *c(H, E)* remains unchanged under permutations of individuals, and of predicates of any family), and axioms that guarantee undogmatic inductive learning, and long-run convergence to relative frequencies. They imply that, for a family $\{P_n\}$, $n = 1, ..., k$, $k > 2$:

$$c_\lambda\text{ (individual } s + 1 \text{ is } P_j,\ s_j \text{ of the first } s \text{ individuals are } P_j) = \frac{s_j + \lambda/k}{s + \lambda},$$

where λ is a positive real number.

The higher the value of λ, the less impact evidence has: induction from what is observed becomes progressively more swamped by a classical-style equal assignment to each of the k possibilities regarding individual $s + 1$.

Significantly, Carnap's various axioms of symmetry are hardly logical truths. More seriously, we cannot impose further symmetry constraints that are seemingly just as plausible as Carnap's, on pain of inconsistency—see Fine (1973, 202). Goodman taught us: that the future will resemble the past in some respect is trivial; that it will resemble the past in all respects is contradictory. And we may continue: that a probability assignment can be made to respect some symmetry is trivial; that one can be made to respect all symmetries is contradictory. This threatens the whole program of logical probability.

14. Frequency interpretations

Frequency interpretations can be thought of as elevating a methodological rule for induction—the straight rule—to the status of a definition of probability. Empiricist in inspiration, and originating with Venn (1876), they identify an event's probability with the *relative frequency* of events of that type within a suitably chosen reference class. The probability that a given coin lands 'heads', for example, might be identified with the relative frequency of 'heads' outcomes in the class of all tosses of that coin. But there is an immediate problem: observed relative frequencies can apparently come apart from true probabilities, as when a fair coin that is tossed ten times happens to land heads every time. Von Mises (1957) offers a more sophisticated formulation based on the notion of a *collective,* rendered precise by Church (1940): a hypothetical infinite sequence of 'attributes' (possible outcomes) of a specified experiment, for which the limiting relative frequency of any attribute exists, and is the same in any recursively specified subsequence. The probability of an attribute A, relative to a collective ω, is then defined as the limiting relative frequency of A in ω. Limiting relative frequencies violate countable additivity (de Finetti 1974).

A notorious problem for any version of frequentism is the so-called *problem of the single case:* we sometimes attribute non-trivial probabilities to results of experiments that occur only once. The move to hypothetical infinite sequences of trials creates its own problems: There is apparently no fact of the matter as to what such a hypothetical sequence would be, nor even what its limiting relative frequency for a given attribute would be, nor indeed whether that limit is even defined; and the limiting relative frequency can be changed to any value one wants by suitably permuting the order of trials. In any case, the empiricist intuition that facts about probabilities are simply facts about patterns in the actual phenomena has been jettisoned.

15. Propensity interpretations

Attempts to locate probabilities 'in the world' are also made by variants of the *propensity* interpretation, championed by such authors as Popper (1959), Mellor (1971) and Giere (1973). Probability is thought of as a physical propensity, or disposition, or tendency of a given type of physical situation to yield an outcome of a certain kind, or to yield a long run relative frequency of such an outcome. This view is explicitly intended to make sense of single-case probabilities. According to Popper, a probability p of an outcome of a certain type is a propensity of a repeatable experiment to produce outcomes of that type with limiting relative frequency p. Given their intimate connection to limiting relative frequencies, such propensities presumably likewise violate countable additivity. Giere explicitly allows single-case propensities, with no mention of frequencies: probability is just a propensity of a repeatable experimental set-up to produce sequences of outcomes. This, however, creates the problem of deriving the desired connection between probabilities and frequencies. This quickly turns into a problem for inductive inference:

it is unclear how frequency information should be brought to bear on hypotheses about propensities that we might entertain.

16. The subjectivist interpretation: Orthodox Bayesianism

Degrees of belief

Subjectivism is the doctrine that probabilities can be regarded as degrees of belief, sometimes called *credences*. It is often called 'Bayesianism' thanks to the important role that Bayes' theorem typically plays in the subjectivist's calculations of probabilities, although this is yet another misnomer since *all* interpretations of probability are equally answerable to the theorem, and subjective probabilities can be defined without any appeal to it. Unlike the logical interpretation (at least as Carnap originally conceived it), subjectivism allows that different agents with the very same evidence can rationally give different probabilities to the same hypothesis.

But what is a degree of belief? A standard analysis invokes betting behaviour: an agent's degree of belief in X is p iff she is prepared to pay up to p units for a bet that pays 1 unit if X, 0 otherwise (de Finetti 1937). It is assumed that she is also prepared to sell that bet for p units. Thus, *opinion* is conceptually tied to certain *behavior*. Critics argue that the two can come apart: an agent may have reason to misrepresent her opinion, or she may not be motivated to act according to her opinion in the way assumed.

Bayesians claim that ideally rational degrees of belief are (at least finitely additive) probabilities. 'Dutch Book' arguments are one line of defense of this claim. A Dutch Book is a series of bets, each of which the agent regards as fair, but which collectively guarantee her loss. De Finetti (1937) proves that *if* your degrees of belief are not finitely additive probabilities, *then* you are susceptible to a Dutch Book. Equally important, and often neglected, is Kemeny's

(1955) converse theorem. A related defense of Bayesianism comes from *utility theory*. Ramsey (1926) and Savage (1954) derive both probabilities and utilities (desirabilities) from preferences constrained by certain putative 'consistency' assumptions.

Updating Probability

Suppose that your degrees of belief are initially represented by a probability function $P_{initial}(–)$, and that you become certain of E (where E is the strongest such proposition). What should be your new probability function P_{new}? The favored updating rule among Bayesians is *conditionalization;* P_{new} is related to $P_{initial}$ as follows:

(Conditionalization) $P_{new}(X) = P_{initial}(X|E)$ (provided $P_{initial}(E) > 0$)

Conditionalization is supported by a 'diachronic' Dutch Book argument (see Lewis 1998): on the assumption that your updating is rule-governed, you are subject to a Dutch book if you do not conditionalize. Equally important is the converse theorem (Skyrms 1987). *Jeffrey conditionalization* allows for less decisive learning experiences in which your probabilities across a partition $\{E_1, E_2,...\}$ change to $\{P_{new}(E_1), P_{new}(E_2), ..., \}$, where none of these values need be 0 or 1:

$$P_{new}(X) = \sum_i P_{initial}(X|E_i)P_{new}(E_i)$$

(Jeffrey 1965). It is again supported by a Dutch book argument (Armendt, 1980). See Diaconis and Zabell (1986) for further probability revision rules.

Orthodox Bayesianism can now be characterized by the following maxims:

B1) Rationality requires an agent's 'prior' (initial) probabilities to conform to the probability calculus.

B2) Rationality requires an agent’s probabilities to update by the rule of (Jeffrey) conditionalization.

B3) Rationality makes no further requirements on an agent’s probabilities.

If orthodox Bayesianism is correct, then there is a sense in which Hume's problem of induction is immediately solved. Inductive inferences based on observational evidence are justified by the appropriate prior subjective probability assignments, suitably updated on that evidence. For example, by B3), rationality permits you to assign:

P_{initial} (the sun will rise on day 10001|the sun rises on days 1, 2, …, 10000) = 0.9999.

Suppose your evidence is:

the sun rises on days 1, 2, …, 10000.

Then conditionalizing on that evidence, as rationally requires according to B2), gives:

P_{new}(the sun will rise on day 10001) = 0.9999.

Similarly, if your prior is of the right form, rationality requires you to assign extremely high probability to all marbles being green after a suitable course of experience with green marbles. And in general, the problem of justifying our inductive practices factors, according to Bayesians, into the problem of justifying the choice of prior, and the problem of justifying conditionalization; and they claim to have made good on both. So far, so good. However, non-Bayesians will find this a Pyrrhic probabilistic victory. For orthodox Bayesianism equally allows priors that would license counterinductive and grue-some inferences, based on the same evidence.

But Bayesianism is a theme that admits of many variations.

17. Unorthodox Bayesianism

Each of B1) – B3) has its opponents. It will prove convenient to revisit them in reverse order.

The suspicion just raised is that orthodox Bayesianism is too permissive: it imposes no constraints on the assignment of priors, besides their conformity to the probability calculus. Rationality, the objection goes, is not so ecumenical. A standard defence (e.g., Savage 1954, Howson and Urbach 1993) appeals to famous 'convergence-to-truth', and 'merger-of-opinion' results. Roughly, their content is that in the long run, the effect of choosing one prior rather than another is attenuated: successive conditionalizations on the evidence will, with probability one, make a given agent eventually converge to the truth, and thus initially discrepant agents eventually come to agreement. In an important sense, at least this much inductive logic is implicit in the probability calculus itself.

Unfortunately, these theorems tell us nothing about how quickly the convergence occurs. In particular, they do not explain the unanimity that we in fact often reach, and often rather rapidly. We will apparently reach the truth 'in the long run'; but as Keynes quipped, "in the long run, we shall all be dead".

Against B3), then, there are more stringent Bayesians who hold this truth to be self-evident: Not all priors are created equal. They thus impose further constraints on priors.

One such constraint is that they be *regular,* or *strictly coherent:* if $P(X) = 1$, then $X = \Omega$ (X is necessary/a logical truth)—see Shimony (1955). It is meant to guard against the sort of dogmatism that no course of learning by (Jeffrey) conditionalization could cure.

We might also want to recognize the role that certain objective facts, or that certain 'expert' opinions, might have in constraining one's subjective probabilities. Call probability function Q an *expert function for* P if the following condition holds:

(*) for all X, $P(X| Q(X) = x) = x$

For example, one might conform one's subjective probabilities to corresponding relative frequencies. With Q being the 'relative frequency' function, (*) becomes a version of the so-called *principle of direct probability*. Or one might think that whatever objective chances might be, they are characterized by their role in conditionally constraining rational credence. With Q being the 'objective chance' function, (*) becomes a version of a principle that, suitably finessed, becomes Lewis' (1980) *Principal Principle*. Or one might argue, as van Fraassen (1995) does, that epistemic integrity requires one ideally to regard one's future opinions as being trustworthy—perhaps because of their having arisen from a rational process of learning. With Q being one's probability function at some future time, (*) becomes a version of van Fraassen's *Reflection Principle*. Q could also encapsulate the opinions of simply *an expert*—a person whom one trusts, for whatever reason.

There have been various proposals for resuscitating symmetry constraints on priors, in the spirit of the classical and logical interpretations. More sophisticated versions of the principle of indifference have been explored by Jaynes (1968). The guiding idea is to maximize the probability function's *entropy*, which for an assignment of positive probabilities $p_1, \ldots, p_n$ to n possibilities equals $-\Sigma_i\, p_i \log(p_i)$.

A set of events (or sentences) is *exchangeable* with respect to a given probability function if every event has the same probability, every conjunction of two events has the same probability, every conjunction of three events has the same probability, and so on. See Skyrms (1994) for an excellent discussion of generalizations of exchangeability, and their use in formulating various Goodmanian theses about projectability. Indeed, commonsense often (but not invariably) seems to require one's probabilities to be exchangeable over 'green'-like hypotheses, but not 'grue'-like hypotheses.

So there are many motivations for rejecting B3). But the suspicion at the end of the last section may still remain: Bayesianism, even with various of these bells and whistles added, is still too permissive. What is wanted is a justification of the 'good' inductive inferences*, and no parallel justification of the 'bad' ones*. These principles do not seem to distinguish the good from the bad. Some of them, on the contrary, only seem to nurture the bad inferences—for example, where we might have hoped to kill off grue-like hypotheses, regularity keeps them all alive. Other principles can play both sides with equal ease: exchangeability, for instance, is characterized purely syntactically, so it can be deployed to vindicate grue-like inferences as well as green-like inferences. Still others, such as the Reflection Principle, seem to be neutral with respect to issues of induction.

B2) also has its opponents. Some authors allow, and even insist upon, other rules for the updating of probabilities besides conditionalization. Jaynes advocates revision to the probability function that maximizes entropy, subject to the relevant constraints.

And some Bayesians drop the requirement that rational probability updating be rule-governed altogether—see van Fraassen (1990a), Earman (1992). Note, however, that in a sense this only makes the problem of induction *worse*. Given that the suggested constraints on the priors do not solve the problem, one might have hoped that the updating rule could take up the slack (and according to the proponents of the convergence results mentioned above, in the long run it does). But if the very requirement of an updating rule is abandoned, then it begins to look as if anything goes: if you want suddenly to jump to a probability distribution that assigns overwhelming probability to all marbles being grue, then you are apparently beyond reproach.

The rejection of B1) is a large topic, and it motivates and can be motivated by some of the non-Kolmogorovian theories of probability, to which we now turn.

18. Non-Kolmogorovian theories of probability

A number of authors would abandon the search for an adequate interpretation of Kolmogorov's probability calculus, since they abandon some part of his axiomatization.

Some authors question its set-theoretic underpinnings. Note that the usual justifications of the probability axioms—Dutch Book arguments and so on—take for granted the sigma-field substructure, rather than justifying it as well. Fine (1973) argues that the requirement that the domain of the probability function be a sigma-field is overly restrictive. Some dispute the requirement that probabilities have numerical values. Fine sympathetically canvasses various theories of comparative probability, exemplified by statements of the form 'A is at least as probable as B'. Then there are advocates of *indeterminate* or of *vague* probabilities, who represent probabilities not as single numbers, but as intervals, or more generally sets of numbers (e.g., Levi 1980; Jeffrey 1983; van Fraassen 1990b). Such vagueness might be vindicated by a set of constraints that go beyond those of the probability calculus, but that fall short of the Carnapian ideal of fixing a unique probability function.

Some dispute the usual constraints on the numerical values. Kolmogorov's probability functions are real-valued. A number of philosophers (e.g., Lewis 1980, Skyrms 1980) allow probabilities to take values from the real numbers of a *nonstandard model* of analysis—see Robinson (1966) or Skyrms (1980) for the construction of such a model. In particular, they allow probabilities to be *infinitesimal:* positive, but smaller than every positive (standard) real number. This can be motivated by a desire to respect both regularity and certain symmetries in infinite probability spaces. Meanwhile, physicists such as Dirac, Wigner, and Feynman have countenanced *negative* probabilities, and Feynman and Cox have flirted with *complex-valued* probabilities. (See Mückenheim (1986) for references.) Renyi (1970) allows probabilities to attain the ‘value’ ∞. We may also want to allow logical/necessary truths to be assigned probability less than one, perhaps to account for the fact that mathematical conjectures may be confirmed to varying degrees—see, e.g. Polya (1968). Thus, mathematics too might be susceptible to induction (to be distinguished from ‘mathematical induction’, a *deductive* argument form!).

Kolmogorov’s most controversial axiom is undoubtedly *continuity*—that is, the ‘infinite part’ of countable additivity. He regarded it as an idealization that finessed the mathematics, but that had no empirical meaning. As we have seen, according to the classical, frequency, and certain propensity interpretations, probabilities violate countable additivity. De Finetti marshals a battery of arguments against it (in the name of subjectivism, but his arguments may be regarded as more general).

Various *non-additive* theories of probability that give up even finite additivity have been proposed—for example, Dempster-Shafer theory, which some regard as codifying the notion of ‘weight of evidence’ (Shafer 1976). So-called “Baconian probabilities” represent another non-additive departure from the probability calculus. The Baconian probability of a conjunction is equal to the minimum of the probabilities of the conjuncts. L.J. Cohen (1970, 1977) regards them as appropriate for measuring inductive support. See Ghirardato (1993) for a survey of non-additive measures of uncertainty, and Howson (1995) for further references.

Lastly, various authors, rather than axiomatizing unconditional probability and defining conditional probability therefrom, take conditional probability as primitive and axiomatize it directly; see Spohn (1986).

19. Some future avenues of research

Having discussed various landmarks of past work in induction and probability, we find ourselves now in the curiously reflexive position of predicting what future work in these areas will look like. Suitably cautioned by the very nature of our subject, and with appropriate degrees of uncertainty, here are some of our best bets.

We think that there is still much research to be done within a broadly Bayesian framework. There are already signs of the rehabilitation of logical probability, and in particular the principle of indifference, by authors such as Stove (1986), Bartha and Johns (forthcoming), Festa (1993) and Maher (2000, forthcoming). This will surely resonate with developments in the theory of infinitesimals, for example within the system of 'surreal numbers' (Conway 1976, Ehrlich forthcoming). Relevant here will also be advances in information theory, randomness and complexity theory (see Fine 1973, Li and Vitanyi 1997), and approaches to statistical model selection, and in particular the 'curve-fitting' problem that attempt to codify simplicity—e.g., the Akaike Information Criterion (see Forster and Sober 1994), the Bayesian Information Criterion (see Kieseppä forthcoming), Minimum Description Length theory (Rissanen 1999) and Minimum Message Length theory (Wallace and Dowe 1999). These may also shed light on the time-honored but all-too-nebulous intuition that 'green'-like hypotheses are somehow 'simpler' than 'grue-like' hypotheses.

Probability theory traditionally presupposes classical set theory/classical logic. There is more work to be done on 'non-classical' probability theory. Bayesians may want to enrich their theory of induction to encompass logical/mathematical learning in response to the so-called 'problem of old evidence' (see Zynda 1995), and to allow for the formulation of new concepts and theories. We also see fertile connections between probability and logic that have been explored under the rubric of 'probabilistic semantics' or 'probability logic'—see Hailperin (1996) and Adams (1998). Roeper and Leblanc (1999) develop such probabilistic semantics for primitive conditional probability functions. More generally we envisage increased attention to the theory of such functions (see, for instance, Festa 1999 for a treatment of Bayesian confirmation theory which takes such functions as primitive, and Hájek forthcoming for general arguments in favor of such functions). Further criteria of adequacy for subjective probabilities will be developed—perhaps refinements of ‘scoring rules’ (Winkler 1996), and more generally, candidates for playing a role for subjective probability analogous to the role that truth plays for belief. There will be more research on the theory of expert functions—for example, in the aggregation of opinions and preferences of multiple experts. This problem is well known to aficionados of the *risk-assessment* literature, which has yet to be mined by philosophers—see Kaplan (1992).

We expect that non-Bayesian research programs will also flourish. Non-additive probabilities are getting impetus from considerations of ‘ambiguity aversion’ (Ghirardato 2001) and ‘plausibility theory’ (Hild 2001). Formal learning theory (see Kelly forthcoming) is also gaining support, and more broadly, philosophers will find much interesting work on induction and learning in the computer science and artificial intelligence literature. And there is a need for more cross-fertilization between Bayesianism and classical statistics, and its recent incarnation in the theory of error statistics (Mayo 1996). For example, hypothesis testing at a constant significance level has long been known to be inconsistent with Bayesian inference and decision theory. Recent work by Schervish, Seidenfeld and Kadane (forthcoming) shows that such ‘incoherence’ is a matter of degree. Moreover, in light of work in the economics literature on ‘bounded rationality’, the study of degrees of incoherence is likely to bear fruit. We foresee related attempts to ‘humanize’ Bayesianism—for example, the further study of vague probability and vague decision theory. And classical statistics, for its part, with its tacit trade-offs between errors and benefits of different kinds, needs to be properly integrated into a more general theory of decision.

Decision theory and the theory of induction will profit from insights in the causal modeling literature. For example, the so-called ‘reference class problem’ arises because a given event-token can typically be placed under indefinitely many event-types; this is what gives the various problems of induction their teeth. But progress can be made when the relevant causes are identified, and techniques along the lines of those developed by Pearl (2000) and Spirtes, Glymour and Scheines (1993) can be appealed to. These techniques will doubtless be finessed. More generally, in this brave new world of inter-disciplinarity and rapid communication, inferential methods developed within one field are increasingly likely to be embraced by practitioners of another.[2]

Bibliography

Adams (1998), *Probability Logic*, CSLI, Stanford University.

Armendt, B. (1980): “Is There a Dutch Book Argument for Probability Kinematics?”, *Philosophy of Science* 47, 583-9.

Arnauld, Antoine (1662): Logic, or, The Art of Thinking ("*The Port Royal Logic")*, tr. J. Dickoff and P. James, Indianapolis: Bobbs-Merrill,1964.

[2] We thank especially Branden Fitelson, Matthias Hild, Chris Hitchcock, Jim Joyce, and Tim Maudlin for helpful comments.

Bartha, Paul and Richard Johns (2001): "Probability and Symmetry", *Philosophy of Science* (Supplemental volume)*;* http://hypatia.ss.uci.edu/lps/psa2k/bartha.pdf.

Bertrand, J. (1889): *Calcul des Probabilités*, 1st edition, Paris: Gauthier-Villars, 2nd edition 1907; reprinted as 3rd edition, New York: Chelsea Publishing Company, 1972.

Carnap, Rudolph (1950): *Logical Foundations of Probability*. University of Chicago Press.

Carnap (1963): “Replies and Systematic Expositions”, in P. A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*, Open Court, La Salle, Ill, 966-998.

Church, A. (1940): “On the Concept of a Random Sequence”, *Bulletin of the American Mathematical Society* 46, 130-135.

Cohen, L. Jonathan (1970): *The Implications of Induction*, London: Methuen.

Cohen, L. Jonathan (1977): *The Probable and the Provable*, Clarendon Press, Oxford.

Conway, John (1976): *On Numbers and Games*, London: Academic Press.

De Finetti, Bruno (1937): “Foresight: Its Logical Laws, Its Subjective Sources”, translated in Kyburg and Smokler (1964), 53 - 118.

De Finetti, Bruno (1974): *Theory of Probability*, John Wiley, New York. Reprinted 1990.

Diaconis, Persi and Sandy Zabell (1986): "Some Alternatives to Bayes's Rule", in *Information Pooling and Group Decision Making,* Proceedings of the Second University of California, Irvine, Conference on Political Economy, Bernard Grofman and Guillermo Owen eds., Greenwich CT: Jai Press Inc, 25-38.

Earman, John (1992): *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, Cambridge, MA: MIT Press.

Ehrlich, Philip (2001): "Number Systems with Simplicity Hierarchies: A Generalization of Conway's Theory of Surreal Numbers," *The Journal of Symbolic Logic*, Vol. 66, No. 3 (September).

Festa, Roberto (1993): *Optimum Inductive Methods : A Study in Inductive Probability, Bayesian Statistics, & Verisimilitude*, Dordrecht: Kluwer Academic Publishers, Synthese Library Series (volume 232).

Festa, Roberto (1999): "Bayesian Confirmation", in M. C. Galavotti and A. Pagnini (eds.), *Experience, Reality, and Scientifc Explanation,* Dordrecht: Kluwer, 55-87.

Fine, Terrence (1973): *Theories of Probability*, New York: Academic Press.

Forster, Malcolm and Elliott Sober (1994): "How to Tell when Simpler, More Unified, or Less Ad Hoc Theories will Provide More Accurate Predictions", *British Journal for the Philosophy of Science* 45: 1-35.

Ghirardato, Paolo (2001): "Coping with Ignorance: Unforeseen Contingencies and Non-Additive Uncertainty". *Economic Theory*, 17:247—276.

Ghirardato, Paolo (1993): "Non-additive Measures of Uncertainty: A Survey of Some Recent Developments in Decision Theory", *Rivista Internazionale di Sciencze Economiche e Commerciali* 40, 253-276.

Giere, R. N. (1973), "Objective Single-Case Probabilities and the Foundations of Statistics", in P. Suppes, L. Henkin, G. C. Moisil, and A. Joja. (eds.), *Logic, Methodology and Philosophy of Science IV*, North Holland, Amsterdam, 467-83.

Good, I. J. (1985): "Weight of Evidence: A Brief Survey", in Bernardo, J. DeGroot, M. Lindley, D. and Smith, A. (eds.) *Bayesian Statistics 2,* Amsterdam: North Holland, 249-269.

Goodman, Nelson (1983): *Fact, Fiction, and Forecast (4th ed.)*, Cambridge, MA: Harvard University Press.

Hájek, Alan (forthcoming): "What Conditional Probability Could Not Be", *Synthese*.

Hailperin, T. (1996): *Sentential Probability Logic*, Bethlehem, PA: Lehigh University Press.

Hempel, Carl G. (1945a): "Studies in the logic of confirmation, Part I", *Mind* 54, 1-26.

Hempel, Carl G. (1945b): "Studies in the logic of confirmation, Part II", *Mind* 54, 97-121.

Hild, Matthias (2001), "Non-Bayesian Rationality", at www.hild.org.

Howson, Colin (1995). "Theories of Probability", *British Journal of Philosophy of Science* 46, 1-32.

Howson, Colin and Peter Urbach (1993). *Scientific Reasoning: The Bayesian Approach*, 2nd ed., Open Court, Chicago.

Hume, David (1739): *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, Oxford, Clarendon Press; 2nd ed.1985

Jaynes, E. T. (1968): "Prior Probabilities", *Institute of Electrical and Electronic Engineers Transactions on Systems Science and Cybernetics*, SSC-4, 227-241.

Jeffrey, Richard (1965): *The Logic of Decision*, University of Chicago; 2nd. ed. 1983.

Jeffrey, Richard (1983): "Bayesianism With a Human Face", "Bayesianism with a Human Face", in *Testing Scientific Theories*, ed. John Earman, Minnesota Studies in the Philosophy of Science, vol. **X,** 133 - 156; reprinted in Jeffrey (1992).

Jeffrey, Richard (1992): *Probability and the Art of Judgment*, Cambridge Studies in Probability, Induction, and Decision Theory, Cambridge: Cambridge University Press.

Jeffreys, Harold (1939): *Theory of Probability*; reprinted in Oxford Classics in the Physical Sciences series, Oxford University Press, 1998.

Johnson, W. E. (1932): “Probability: The Deductive and Inductive Problems”, *Mind* 49, 409-423.

Kahneman, D., P. Slovic and A. Tversky (eds.) (1982): *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.

Kaplan, Stan (1992): "'Expert Information' Versus 'Expert Opinions'. Another Approach to the Problem of Eliciting/Combining/Using Expert Knowledge in PRA", *Reliability Engineering and System Safety* 35, 61-72.

Kelly, K, (1996): *The Logic of Reliable Inquiry*, Oxford: Oxford University Press.

Kemeny, J. G. (1955): “Fair Bets and Inductive Probabilities”, *Journal of Symbolic Logic* 20, 263-273.

Keynes, J. M. (1921): *Treatise on Probability*, Macmillan, London. Reprinted 1962, Harper and Row, New York.

Kieseppä, I. A. (2001): "Statistical Model Selection Criteria and Bayesianism", *Philosophy of Science*(Supplemental volume).

Kolmogorov, A. N. (1933): *Grundbegriffe der Wahrscheinlichkeitrechnung*, Ergebnisse Der Mathematik; translated as *Foundations of Probability*, New York: Chelsea Publishing Company, 1950.

Kornblith, Hilary (1985): *Naturalizing Epistemology,* Cambridge, MA: MIT Press.

Kyburg, Henry E. Jr. and Howard E. Smokler (eds.) (1964): *Studies in Subjective Probability*, New York: John Wiley and Sons. 2nd edition 1980.

Laplace, Pierre Simon de (1814): “Essai Philosophique sur les Probabilités”, Paris. Translated into English as *A Philosophical Essay on Probabilities*, New York, 1952.

Levi, Isaac (1980): *The Enterprise of Knowledge,* MIT Press, Cambridge, Massachussetts.

Lewis, David (1980): “A Subjectivist's Guide to Objective Chance”, in *Studies in Inductive Logic and Probability*, Vol II., University of California Press, 263-293; reprinted in Lewis 1986.

Lewis, David (1986): *Philosophical Papers,* Volume II, Oxford: Oxford University Press.

Lewis, David (1998): *Papers in Philosophical Logic*, Cambridge: Cambridge University Press.

Ming Li and Paul Vitanyi (1997): *An Introduction to Kolmogorov Complexity and its Applications*, 2nd edition, Springer-Verlag, New York.

Maher, Patrick (2000): "Probabilities for Two Properties", *Erkenntnis* 52, 63-91.

Maher, Patrick (forthcoming): "Probabilities for Multiple Properties: The Models of Hesse and Carnap and Kemeny", *Erkenntnis*.

Mayo, Deborah G. (1996): *Error and the Growth of Experimental Knowledge*, University of Chicago Press.

Mellor, D.H. (1971): *The Matter of Chance*, Cambridge University Press, Cambridge.

Milne, Peter (1996): “Log[P(h/eb)/P(h/b)] is the one true measure of confirmation”, *Philosophy of Science* 63, 21-26.

Mückenheim, W. (1986): "A Review of Extended Probabilities", *Physics Reports*, Vol. 133, No. 6, 337-401.

Pearl, Judea (2000): *Causality*, Cambridge: Cambridge University Press.

Polya, G. (1968): *Patterns of Plausible Inference*, Princeton University Press, Princeton, 2nd ed.

Popper, Karl (1959): “The Propensity Interpretation of Probability, *British Journal for the Philosophy of Science* 10, 25-42.

Popper, Karl R. (1968): *The Logic of Scientific Discovery*, London: Hutchinson & Co, revised edition.

Putnam, Hilary (1974): “The ‘Corroboration’ of Theories”, reprinted in his *Mathematics, Matter, and Method*, Cambridge University Press 1979, 250-69.

Ramsey, F. P. (1926): “Truth and Probability”, in *Foundations of Mathematics and other Essays;* reprinted in Kyburg and Smokler (1964), and in D.H. Mellor (ed.) *Philosophical Papers*, Cambridge University Press, Cambridge, 1990.

Reichenbach, Hans (1938): *Experience and Prediction,* Chicago: University of Chicago Press.

Reichenbach, Hans (1949): *The Theory of Probability,* Berkeley: University of California Press.

Renyi, Alfred (1970): *Foundations of Probability*, San Francisco: Holden-Day, Inc.

Rissanen, J. (1999): "Hypothesis selection and testing by the MDL principle", *Computer Journal* 42 (4), 260-269.

Robinson, Abraham (1966): *Non-Standard Analysis*, North Holland Publishing Company, Amsterdam.

Roeper, P. and Leblanc, H. (1999): *Probability Theory and Probability Logic*, Toronto Studies in Philosophy, University of Toronto Press.

Salmon, Wesley (1967): *The Foundations of Scientific Inference,* Univ. of Pittsburgh Press.

Savage, L. J. (1954): *The Foundations of Statistics*, New York, John Wiley.

Schervish, Mark, J, Teddy Seidenfeld and Joseph B. Kadane (2002): "How Incoherent is Fixed-Level Testing?", *Philosophy of Science* (Supplemental volume).

http://hypatia.ss.uci.edu/lps/psa2k/incoherent-fixed-level.pdf.

Shafer, Glenn (1976): *A Mathematical Theory of Evidence*, Princeton University Press, Princeton.

Shimony, Abner (1955): "Coherence and the Axioms of Confirmation", *Journal of Symbolic Logic* 20, 1-28.

Skyrms, Brian (1980): *Causal Necessity*, Yale University Press.

Skyrms, Brian (1987): "Dynamic Coherence and Probability Kinematics", *Philosophy of Science* 54, 1-20.

Skyrms, Brian (1994): "Bayesian Projectibility", in *Grue!* (D. Stalker ed.), Chicago: Open Court, 241-62.

Spirtes, P., C. Glymour and R. Scheines (1993): Causation, Prediction, and Search, New York: Springer-Verlag.

Spohn, Wolfgang (1986): "The Representation of Popper Measures", *Topoi* 5, 69-74.

Stove, D. C. (1986): *The Rationality of Induction,* Oxford: Oxford University Press.

Strawson, Peter (1952): *Introduction to Logical Theory*, New York: John Wiley & Sons.

van Cleve, James (1984): "Reliability, Justification, and the Problem of Induction", *Midwest Studies in Philosophy* 1984, 555-568.

van Fraassen, Bas (1990a): "Rationality Does Not Require Conditionalization", in E. Ullman-Margalit (ed.), *The Israel Colloquium: Studies in History, Philosophy and Sociology of Science*, Kluwer, Dordrecht.

van Fraassen, Bas (1990b): "Figures in a Probability Landscape", in J.M. Dunn and A. Gupta (eds.), *Truth or Consequences*, Dordrecht: Kluwer, 345-356.

van Fraassen, Bas (1995): "Belief and the Problem of Ulysses and the Sirens", *Philosophical Studies*. **77**, 7-37.

Venn, John (1876): *The Logic of Chance*, 2nd ed., London: Macmillan and co.

von Mises, R (1957): *Probability, Statistics and Truth*, 2nd edition, revised English edition, New York: Macmillan.

Wallace, C.S. and D.L. Dowe (1999): "Minimum Message Length and Kolmogorov Complexity", *Computer Journal* (special issue on Kolmogorov complexity), Vol. 42, No. 4, 270-283.

Winkler, R. L. (1996): "Scoring Rules and the Evaluation of Probabilities", *Test,* Vol. 5, No. 1, 1-60.

Zynda, Lyle (1995): "Old Evidence and New Theories," *Philosophical Studies* 77, pp. 67-95.