Part XI

INDUCTIVE, FUZZY, AND QUANTUM PROBABILITY LOGICS

35

Inductive Logic

STEPHEN GLAISTER

All inductive logicians aim to construct a formally articulated theory of good ampliative (non-deductive) inference that parallels existing formal theories of good deductive inference. They disagree, however, about the extent and respects of that parallel, as well as about the exact formal resources that should be brought to bear. What I will call *Good Old-Fashioned Inductive Logic* (GOFIL) holds that the parallels between deductive logic and inductive logic are straightforward and extensive. ('GOFIL' and related acronyms follow a well-known model due to John Haugeland.) On this view inductive logic, like deductive logic, studies arguments, but whereas deductive logic studies the relation of deductive validity between an argument's premises and its conclusion, inductive logic studies the degree to which those premises support or confirm that conclusion.

The obvious instrumentality with which to articulate this system of degrees is probability. The basic properties of probability are codified by axiomatizations such as those of Kolmogorov and of Renyi. Advocates of GOFIL, together with many statisticians and essentially all philosophers of probability, hold that it makes perfectly good sense to ask – even that it is essential we ask – what probability *is* beyond those basic formal properties (Salmon 1967; Walley 1991; Hájek 1997). Taking existing uses of probability concepts in both commonsense and science as the first word in matters of extra-formal interpretation, GOFIL suggests that *one* thing probability is, particularly in epistemic contexts, is a parameter expressing degree of confirmation. That is, the probability $P(B \mid A)$ that conclusion proposition B is true given premise proposition A ($= \cap A_i$ if the argument has multiple premises) is understood as a measure of the objective, logical degree to which A supports or confirms B. GOFIL therefore sponsors a so-called 'logical' interpretation of (two-place) probability (Keynes 1921; Jeffreys 1957; Carnap 1962). Since Carnap's version of GOFIL is the most developed and influential we will concentrate on that account. In section 1, then, we review the principal achievements of and challenges faced by GOFIL á là Carnap.

Many of GOFIL's achievements are detachable from that program's commitment to a logical interpretation of probability. In section 2, we survey the reincarnation of GOFIL in the context of a subjectivist interpretation of probability: a development we will call *Subjectivist Inductive Logic* (SIL).

SIL rejects GOFIL's logical account of probability, but it largely perpetuates GOFIL's basic conception of inductive logic as a matter of articulating standards of coherence and consistency that can be used to assess particular inferences or inference forms. A more radical departure from GOFIL refocuses attention away from issues of coherence and consistency, and towards the study of various sorts of logical guarantees of convergence to the truth. We discuss this *New-Fangled Inductive Logic* (NFIL) in section 3.

1 Good Old-Fashioned Inductive Logic (GOFIL): Carnap's Program

Carnap developed his account of inductive logic through a long series of important publications over 30 years, reaching an apex of both generality and compatibility with standard probabilistic terminology in the posthumously published Carnap (1971, 1980). In this section, we will employ essentially the terminology used in this later work.

Formal preliminaries

Let a *family of properties*, $\{F_1, \ldots, F_k\}$ be a set of properties that are pairwise exclusive and jointly exhaustive. Each property in the family thus functions as a complete characterization of an individual in the logic. An *atomic proposition* is a proposition that ascribes one of the properties, F_i, to an individual, for example F_3b. A *sample proposition* is a finite conjunction of atomic propositions in which each atomic proposition involves a different individual, for example $F_3b \cap F_3c \cap F_4d$. *n(E)* or n (if the reference is clear) is the total number of individuals involved in E, and *ind(E)* is the set of individuals involved in E. Let the empty sample proposition be the necessarily true proposition, Ω.

$n_i(E)$ or just n_i (if the reference is clear) is the number of individuals to which E ascribes property F_i, and $\mathbf{n}(E)$ or $\mathbf{n}$ (if the reference is clear) is the frequency vector $\langle n_1, \ldots, n_k\rangle$ for E. The number of possible frequency vectors for a sample proposition E is $\binom{n+k-1}{k}$, where $n = n(E) = \Sigma n_i$. $\left(\text{Note that } \binom{n}{r} = \frac{n!}{r!(n-r)!}.\right)$ The number of possible sample propositions involving exactly ind(E) with a given frequency vector $\mathbf{n}$ is given by the multinomial coefficient for that vector, $\frac{n!}{n_1!n_2!\ldots n_k!}$ If there are q individuals overall then there are $\binom{q}{n}$ distinct sets of n individuals available to be partitioned by $\mathbf{n}$, hence $\binom{q}{n} \cdot \frac{n!}{n_1!n_2!\ldots n_k!}$ possible realizations of $\mathbf{n}$ in that population.

Let P be a probability function on the algebra of atomic propositions for countably many individuals. A *singular predictive inference* is a conditional probability, $P(F_ia \mid E)$, where E is a sample proposition that does not involve individual a. A *rule of succession* is a general formula for P's singular predictive inferences. We define the unconditional probability $P(F_ia)$ as $P(F_ia \mid E = \Omega)$. Unconditional probabilities for arbitrary sample propositions follow immediately, for example:

$$P(F_ia \cap F_jb) = P(F_ia) \cdot P(F_jb \mid F_ia)$$
$$P(F_ia \cap F_jb \cap F_kc) = P(F_ia) \cdot P(F_jb \mid F_ia) \cdot P(F_kc \mid F_ia \cap F_jb)$$

and so on. All other unconditional probabilities follow by additivity of the probabilities of these basic possibilities.

P is *exchangeable* just in case the probabilities of sample propositions are functions of their frequency vectors, that is P(E) = P(E′) if **n**(E) = **n**(E′). Put another way, P is exchangeable just in case one doesn't change probabilities merely by altering which individuals have which properties.

To the continuum and beyond

Carnap's first major work in inductive logic (Carnap 1962) culminates in a long appendix on a particular logical probability or degree of confirmation function, c*. The rule of succession for c* is

$$c^*(F_i a \mid E) = \frac{n_i + 1}{n + k}$$

that is where E involves n individuals of which n_i have property F_i.

C* gives equal prior probabilities for an individual having an arbitrary property, that is for all F_i, $c^*(F_i a \mid \Omega) = 1/k$, since $n_i = n = 0$. This implies both that each possible frequency vector for the whole population is allotted the same prior probability, and that that allotment is split evenly among its realizations. For example, when the universe of individuals comprises two coin tosses, a and b, and the family of properties is just $\{F_1 = \text{‘heads’}, F_2 = \text{‘tails’}\}$ then the c* priors are:

$$c^*(F_1 a \cap F_1 b \mid \Omega) = c^*(F_2 a \cap F_2 b \mid \Omega) = 1/3$$
$$c^*(F_1 a \cap F_2 b \mid \Omega) = c^*(F_2 a \cap F_1 b \mid \Omega) = 1/6$$

That is, the outcomes of 0, 1, and 2 heads are given equal weight, notwithstanding that there are more ways to get exactly 1 head.

Cases such as this almost immediately start one worrying that an alternative to c* that assigns equal prior probability (here 1/4) directly to all the possible realizations of all frequency vectors might be preferable. Carnap (1962) called this alternative, $c^\dagger$, and noted that Peirce, Keynes, and Wittgenstein had all succumbed to its charms. In a tour de force, however, Carnap showed that the rule of succession for $c^\dagger$ is:

$$c^\dagger(F_i a \mid E) = \frac{1}{k}$$

That is, $c^\dagger$ makes a certain sort of empirical learning impossible: its singular predictive inferences ignore all *observed frequency* information.

Whence comes the appeal of $c^\dagger$, say in the coin-tossing case above, if it is, in the abstract, inductively catastrophic? Evidently its appeal in the case at hand is grounded in the fact that we are jumping to the conclusion that the coin to be tossed is (close to) objectively fair. If we assume this (or indeed any other particular bias for the coin) then there is an important sense, underlined by Carnap's result for $c^\dagger$, in which we already know everything there is to know about the coin we are tossing. The exact sequence of heads and tails remains to be determined, of course, but that's just the unfolding of a chance process: it's 'whatever happens.' And if we know the chance parameters for the overall process then no stage of the unfolding chance process tells us anything about any other stage of that process. $c^{\dagger}$ is not an inductive catastrophe in the case of tossing a coin with known bias, rather it's a legitimate expression of the fact that there's nothing left to learn about the case at hand.

Reflecting on $c^{\dagger}$ in this way helps us see that the basic inductive problem for Carnap is equivalent (given two outcomes) to trying to figure out the bias of a coin from the actual outcomes of a series of tosses. From this perspective, the degeneracy of $c^{\dagger}$ is just that it is appropriate only for a case in which exactly *that* inductive problem has already been solved. In actual empirical applications, moreover, we are always open to revisiting our estimates of a coin's bias – a string of 1000 heads from a coin we believed fair would *always* give us pause. It follows that $c^{\dagger}$ is at best a contextually specific, convenient approximation; one which Carnap himself analogizes to the use of 22/7 to approximate π.

At the opposite extreme from $c^{\dagger}$ is the so-called 'straight rule,' c^{sr}. This alternative to c^* ignores all *prior probability* information and simply predicts the continuation of the components of the sample proposition's frequency vector into the future, that is:

$$c^{sr}(F_i a \mid E) = \frac{n_i}{n}.$$

We can think of c^{sr} as at one end of a spectrum, giving no weight to the prior probabilities, $c^{\dagger}$ at the other end, giving incomparably great weight to the priors, and c^* as somewhere in between. Carnap (1952) makes the obvious weight parameter explicit, yielding the following family of rules of succession:

$$c^{\lambda}(F_i a \mid E) = \frac{n_i + \lambda/k}{n+\lambda}$$

where $0 \leq \lambda \leq \infty$. Since c^{λ} has continuum many instances, Carnap called this system the 'Continuum of Inductive Methods.' C^{λ} reduces to c^{sr}, c^*, and $c^{\dagger}$ when λ is 0, k, and ∞ respectively.

Carnap (1980) generalizes still further, dropping the requirement of uniformity of the prior probability, γ_i, that an individual will have property F_i:

$$c^{\lambda,\gamma}(F_i a \mid E) = \frac{n_i + \lambda\gamma_i}{n+\lambda}$$

where λ is positive and finite and $\Sigma\gamma_i = 1$. (Carnap (1980) analyzes extreme rules such as c^{sr} and $c^{\dagger}$ only in the limit, as λ approaches 0 and ∞ respectively.) We can think of the probabilities generated by this sort of rule as a matter of first augmenting the n-membered sample population with a virtual population comprising λ individuals with frequency vector $\langle \lambda\gamma_1, \lambda\gamma_2, \ldots, \lambda\gamma_k \rangle$, then recalculating the relative frequencies from there. C^* is the case where each property gets exactly one virtual representative (Jeffrey 1980: 2–3).

Carnap presented many different sets of *qualitative* conditions on P over the years, each of which he intended to be sufficient for his favored family of logical probabilities at the time. Carnap (1980: section 19) proves that the following relatively sparse group of conditions is sufficient for the λ–γ-continuum of inductive methods:

(λ–γ1) P is exchangeable.
(λ–γ2) P is regular: $P(E) > 0$ for all E.
(λ–γ3) Strict Instantial Relevance: $P(F_ib \mid F_ia) > P(F_ib)$.
(λ–γ4) Sufficientness: $P(F_ia \mid E)$ is a function just of n and n_i.

Note that $c^{sr} = c^{0,\gamma}$ conflicts with (λ–γ2) since if $n_i(E) = 0$ then $c^{0,\gamma}(E) = 0$, and $c^{\infty,\gamma}$ conflicts with (λ–γ3) since $c^{\infty,\gamma}(F_ib \mid F_ia) = c^{\infty,\gamma}(F_ib)$. Lastly, note that ($\lambda$–$\gamma$4) is vacuous if there are only two possible properties, so that, strictly speaking, the given postulates only imply the λ–γ-continuum for $k \geq 3$. Carnap saw the problem, and solved it – inelegantly – by adducing a *quantitative* axiom of linearity to cover the $k = 2$ case. We set aside this unfortunate wrinkle in Carnap's approach here.

The basic problem

Before discussing relatively technical objections to and further developments of Carnap's program, it is worth asking about the extent to which that program succeeds in meeting its original goals. Recall that the basic suggestion of GOFIL was that *one* thing probability could be, particularly in epistemic contexts, is a parameter registering degree of confirmation.

Now if, say, c* had emerged as a uniquely compelling inductive method then it would have been possible for Carnap to declare victory: to say that probability in many epistemic contexts just *is* c*. But if, as Carnap clearly believed by 1952, c* is *not* uniquely compelling so that, so to speak, degrees of confirmation are *many* while probabilities are *one*, then the logical interpretation of probability has to be abandoned. $P(B \mid A)$ *can't* just be the degree to which A confirms B since degree of confirmation turns out to have additional argument places that probability lacks. Moreover, Carnap agrees that the value of λ "is fundamentally not a theoretical question" but a "practical decision," albeit one which can be importantly informed by "theoretical results concerning the properties of the various inductive methods" (Carnap 1952: 53). The λ–γ-continuum, of course, only expands the role of practical decision making. But then, whatever else he might have done, Carnap hasn't provided a logical interpretation of (two-place) probabilities, fixed as relations of entailment are fixed, simply by the nature of the underlying field of propositions. And assigning all particular degrees of support only relative to λ- and γ-values that are themselves matters of decision raises the specter of circularity, or at least of a kind of holism about probability (i.e. if those decisions, as it's natural to suppose, already involve probabilistic reasoning of some kind). This is unpromising ground on which to try to erect a strictly logical interpretation of probability.

From the late 1950s onward, and especially in Carnap (1963), Carnap strongly emphasizes the role of degrees of confirmation in helping determine expected utilities, fair betting quotients, and so on. It is not clear whether Carnap intended this new emphasis to solve or to concede defeat by the problems raised in this section. Whatever

Carnap may have intended, strongly emphasizing betting behavior and decision making *invites* exploration of how much of GOFIL can be retained given a subjectivist interpretation of probability. And historically this path has been very popular. Indeed, in one obvious respect, GOFIL apparatus immediately acquires a new luster in a subjectivist setting: symmetry arguments and principles that are endlessly controversial when wielded as additional universal postulates to help fix logical probabilities, are necessarily less objectionable when employed opportunistically, as tools for forming subjective probability models of particular cases. We consider subjectivized inductive logic in detail in section 2.

Other problems and developments

In this section, we briefly review some relatively technical problems for GOFIL.

Confirming universal generalizations

In domains where there are infinitely many individuals all of Carnap's inductive methods give zero prior probabilities (hence – except for c^{sr} – also zero posterior probabilities on finite evidence) to universal generalizations (UGs). One response is simply to accept the consequence, fashioning the point either as a sobering reminder of how far the literally universal outstrips our abilities to probabilize (R. Price, De Morgan, Jeffrey), or as a demonstration of how far the mathematics of infinity takes us away from the sorts of epistemic contexts that matter (Ramsey, Savage, T. Fine). The other main response is, of course, to try to modify Carnap's apparatus to permit assigning positive prior probability to UGs in infinite domains (and swifter confirmation in finite domains). From a subjectivist perspective, the problem is no sooner stated than it is solved: simply put finite probability where it's needed and make appropriate adjustments elsewhere (Jeffreys and Wrinch 1919; Earman 1992: 89–90). But how to justify this sort of flexibility within a GOFIL setting?

Two attempts have been made to meet this challenge, both of which centrally involve amending (λ–γ4) (actually both amendments address only the λ-continuum, but the difference doesn't matter here). Zabell (1997b) proposes the following minimal modification:

(λ–γ4.1) $P(P_i a \mid E)$ is a function just of n, n_i, *except* when E involves only a single property.

and proves a remarkable theorem showing that essentially just exchangeability of P determines both:

- the existence of prior and posterior probabilities for each UG, $(\forall j)F_i a_j$;
- a formula that makes $P(F_i a \mid E)$ in these cases a weighted average of its λ-continuum value and the posterior probability for $(\forall j)F_i a_j$.

The most extended response to the problem of confirming UGs within GOFIL is due to Hintikka and Niiniluoto (1980). The core of this response is the weakening of (λ–γ4) to:

(λ–γ4.2) $P(P_ia \mid E)$ is a function just of n, n_i, and of the number of distinct properties *not* involved in E.

The thought behind the additional argument parameter here is that "it determines the number of nonequivalent generalizations compatible with the sample" (Hintikka and Niiniluoto 1980: 160). Coordinately, the underlying technical innovation of the so-called H-N systems is to assign probabilities in the first instance directly to generalizations about which properties (and relations) are instantiated in a population (the 'constituents' of Hintikka (1966)). One can, it turns out, do this in a way that permits (1) the calculation of all sample proposition probabilities, (2) positive probability for UGs independently of the cardinality of the domain, and (3) much faster confirmation of generalizations in finite domains. Carnap's λ-continuum even emerges as the sole H-N system in which UGs fail to receive positive probability in infinite domains (Hintikka and Niiniluoto 1980: 173).

New properties/species

Carnap's inductive methods (as well as the H-N systems) suppose that we know all the basic properties, $\{F_i\}$, in advance. But this is deeply unrealistic: real-world inductions involve learning about new types almost as much as they do learning about new tokens (of pre-digested types). Zabell (1992, 1997a) shows how to make a Carnapian framework more realistic, by allowing for singular predictive inferences about novel properties or species.

Suppose for ease of exposition that the observables are letters in the alphabet. Evidently the frequency vector (from n_A to n_Z) for the sample of observations DFBBBAA, $\langle 2, 3, 0, 1, 0, 1, 0, 0, \ldots, 0\rangle$ is not an appropriate statistic unless we possess a prior enumeration of all 26 types. If we knew *just* DFBBBAA, then the only available frequency vector would seem to be the vector comprising the frequencies of the actually observed types (from n_A to n_F), $\langle 2, 3, 1, 1\rangle$. We can further imagine interpreting the vector in the following minimal fashion:

> "fourth species observed showed up twice," "third species observed showed up three times,". . . .

If we now suppose that possible observations are indexed by times $\{1, \ldots, n\}$, then the frequency vector just for observed types can be thought of as constituting a partition π of the index set $\{1, \ldots, 7\}$. We can now generate a higher-order statistic for a partition corresponding to the *frequencies of the frequencies* in the partition of the index set. Let a_j be the number of types with j observed tokens, that is the number of j-membered partition cells in the partition of the index set. For example, a_1(DFBBBAA) = 2, a_2(DFBBBAA) = a_3(DFBBBAA) = 1.

Let Π_n be a random variable taking as values possible partitions of $\{1, \ldots, n\}$ and let $\mathbf{a} = \langle a_1, \ldots, a_n\rangle$ be the *partition vector* for an n-membered sample, so that, for example, $\mathbf{a}$(DFBBBAA) = $\langle 2, 1, 1, 0, 0, 0, 0\rangle$. Suppose finally that one has a probability over the space of possible partitions of $\{1, \ldots, n\}$ and say that that probability is *partition exchangeable* iff all partitions with the same partition vector are equiprobable, that is $P(\Pi_n = \pi_1) = P(\Pi_n = \pi_2)$ if $a(\pi_1) = a(\pi_2)$, where π_1 and π_2 are partitions of $\{1, \ldots, n\}$.

Zabell (1997a) shows that a new, three-parametered continuum of inductive methods is implied by partition exchangeability, and three other conditions. The first two conditions are just partition counterparts of regularity (λ–γ2) and sufficientness (λ–γ4). The final condition governs the probability that the next individual is of a novel species:

(Z) $P(e_{n+1} \in S_{t+1} \mid \langle n_1, \ldots, n_t \rangle)$ is a function just of the number of species already observed, t, and the sample size, n.

See Zabell (1997a: section 2) for the new continuum of predictive probabilities itself. For some of the relations between the new continuum and the H-N systems, see Zabell (1992: 218).

Note finally that some authors (Salmon 1967; Fine 1973) have worried that Carnapian degree of confirmation values are inappropriately sensitive to refinements in the space of properties. Zabell's work, however, appears to demonstrate one way in which such sensitivity is not only appropriate but essential.

Analogy

Carnap's inductive logic can be understood as importantly anti-analogical. Whereas exchangeability implies that order of individuals is unimportant, proximity of individuals in some ordering or metric is often taken to be a reasonable basis for inference. Similarly, whereas sufficientness conditions insulate predictions about one type from (frequency) information about other types, proximity of types in some ordering or metric is often taken to be a reasonable basis for inferring features of one type from another.

We briefly discuss the weakenings of exchangeability that are needed for models of the first sort of analogical reasoning in the discussion of "De Finetti's exchangeability reduction" below. Notable attempts to model the second sort of analogy in a broadly Carnapian spirit include Niiniluoto (1981), Constantini (1983), Kuipers (1984), Skyrms (1993), and Maher (2000).

2 Subjectivized Inductive Logic (SIL): De Finetti Regnant

We observed in the section "To the continuum and beyond", above, that Carnap thinks of the basic inductive problem as analogous to trying to divine the bias of a coin from an actual sequence of tosses. The crucial point is that the system being theorized about is supposed, conditional on any particular bias value(s), to produce sequences of (objectively) independent, identically distributed (same objective probabilities each time for the different possible outcomes) trials.

This situation is one of the most well-studied problems in statistics, and particularly in Bayesian statistics. From a Bayesian statistical perspective, the problem is just to choose an appropriate prior probability on the possible values of the bias parameter so that (1) computation of posteriors is easy, and (2) convergence to the true bias is guaranteed. Sometimes statisticians recommend a uniform or 'flat' prior for these purposes. At least equally commonly, however, they allow any member of the family of priors that essentially share the functional forms of the likelihoods (i.e. the probabilities of sample propositions given values of the relevant bias parameter(s)) as functions of the bias parameter(s). Most distinguished are the so-called *natural conjugate* priors. (When the product of the prior and the likelihood yields a posterior distribution in the same family as the prior, the prior is said to *conjugate* with the likelihood function. When that prior conjugates with that likelihood by essentially sharing its functional form then the prior is *natural*.) If the system generating the sample is binomial, the natural conjugate priors are the Beta distributions; if it is multinomial the natural conjugate priors are the Dirichlet distributions. These distributions themselves have multiple parameters. If these parameters have uniform values the resulting Beta and Dirichlet distributions are said to be symmetric. Flat priors result if that uniform value is 1 (Festa 1993: chapter 6; Tanner 1996).

Remarkably, the flat prior corresponds exactly to Carnap's c*, the *symmetric* natural conjugate priors to the λ-continuum, and the natural conjugate priors *in toto* to the λ–γ-continuum. GOFIL subjectivized – SIL – just is Bayesian statistics. The great connecting principle here is Carnap's requirement that degree of confirmation functions be exchangeable, that is (λ–γ1). Famously, De Finetti (1937) proved that any infinite sequence of random variables (i.e. one for each trial) for which every finite subsequence is exchangeable (i.e. according to a subjective probability P over those infinite sequences of trials) has a unique representation as a (possibly continuously) weighted average or mixture of probabilities, each one of which makes the random variables (r.v.s) independent and identically distributed (IID). We will state De Finetti's result precisely just for the binomial case:

De Finetti Representation Theorem Let $\{X_i\}_{i=1}^{\infty}$ be an infinite sequence of {0, 1}-valued random variables with $\{X_i\}_{i=1}^{n}$ exchangeable for each n (according to P); then there is a unique probability measure μ on [0,1] such that for each fixed sequence of zeros and ones $\{e_i\}_{i=1}^{n}$ we have

$$P(X_1 = e_1, \ldots, X_n = e_n) = \int_0^1 p^k (1-p)^{n-k} \, d\mu(p)$$

where

$$k = \sum_{i=1}^{n} e_i \text{ (i.e. k is the number of 'successes')}.$$

This theorem, together with its relatives for sequences of multinomial r.v.s, real-valued r.v.s (De Finetti 1937), and beyond (Hewitt and Savage 1955), implies that choosing c*, c^{λ}, and $c^{\lambda,\gamma}$ as rules of succession is equivalent to choosing the various distributions (or distribution families) mentioned above as mixing measures for the relevant version of De Finetti's Theorem.

Probabilities with respect to infinite sequences of random variables are sometimes decried as unrealistic (Jeffrey 1992), and in part for this reason, finite exchangeable sequences of r.v.s have also been extensively studied. Probabilities for such sequences have unique representations as mixtures of (non-IID) hypergeometric sequences. Extendibility of a finite exchangeable sequence to longer and longer finite exchangeable sequences, however, ensures convergence to representability by a mixture of IID sequences (Diaconis 1977). We set aside the case of finite exchangeability here.

Conditionalizing an exchangeable probability P on outcomes of trials leaves the subsequences of remaining outcomes exchangeable. By Bayes' theorem and the weak law of large numbers, the weights of subsequent mixing measures gradually become focused on a single IID sequence, corresponding to a single parameter value (or vector of parameter values) *unless* one's original mixing measure starts out strangely skewed away from the true parameters of the system, a possibility that natural conjugacy blocks (Diaconis and Freedman 1986). The power of De Finetti's representation theorem is that it shows how this elegant model of learning from experience is implicit in little more than an assumption of a particular sort of subjective indifference or symmetry in one's personal probabilities – exchangeability – together with the assumption that new information is assimilated via conditionalization. As we saw above, De Finetti's theorem also clarifies how to understand Carnap's efforts from a subjectivist standpoint. Over and above identifying the exchangeable probabilities, Carnap's various conditions can be seen as limning the properties of various families of prior mixing measures. De Finetti's result also suggests a more general perspective, according to which to equip a subjective probability with a symmetry of some kind just *is* to endow the agent in question with a conception of objective chance. This vision, which can be pursued through more and more abstract symmetries, often with mathematical roots independent of De Finetti's work, for example in ergodic theory, has proved tantalizing (Skyrms 1984: chapter 3, 1994).

De Finetti himself boldly made two further claims on behalf of his theorem (and its supporting materials): that it paved the way for the complete elimination of objective probability or chance parameters from statistics, and that it solved Hume's problem of induction. Let us briefly consider these claims in turn.

De Finetti's exchangeability reduction

An immediate, technical obstacle to any *general* reduction of IID notions – objective independence and objective equiprobability – to exchangeability is that if the infinite sequence of random variables take values in very rich spaces then no representation in terms of mixtures of sequences of IID trials for those r.v.s may be possible (Dubins and Freedman 1979). But let us set aside this relatively technical worry here.

In many cases, one needs to construct a subjective probability for a situation or phenomenon. And when the phenomenon is an infinite (or infinitely extendible) sequence of random variables it makes sense to ask whether exchangeability or some other related symmetry assumption is justified or reasonable. It certainly *looks* as though the better part of that justification will be an appeal to background knowledge about the phenomenon in question, to our understanding of how 'coin toss' – or 'urn model' – like the phenomenon is. If the phenomenon is judged to be 'coin toss' – like – or in the simplest case just *is* the tossing of a coin of some kind, then an IID sequence can be reasonably expected, and, in effect, only the constant probability of success parameter remains to be determined. But if we had some specific and contrary background knowledge about the coin in question, for example if we knew that the coin was made of some highly unstable material such that every heads outcome increases the probability of heads on the next toss, then it would be perverse to assume exchangeability. In this sort of case the *order* of outcomes matters and not just the frequency vector, hence exchangeability is inappropriate.

Subjectivists have developed models of phenomena that objectivists would describe as exhibiting various sorts of parameterized dependency, under the generic heading of *partial exchangeability* (De Finetti 1938). The best explored of these is Markov exchangeability which focuses not on invariance of probabilities under permutation of trials (and on frequency vectors) but on invariance under switching of sub-sequences of trials that share starting and ending points (and on vectors of initial states and transition counts). See Diaconis and Freedman (1980) and Skyrms (1994: section 5) for further discussion.

The work in this area is impressive, and constitutes an absolutely essential broadening of the base for De Finetti's reductive proposal. It seems unlikely, however, that it does much more than push our basic objection back a step. Even with a wider arrange of symmetries to appeal to, the subjectivist still seems to have to play catch-up with respect to the objectivist. There are, after all, essentially unlimited forms of dependency and objective eccentric character, so that it is hard to see how to avoid the conclusion that symmetries in subjective probabilities are normally best seen as responses to background knowledge about objective symmetry and dependency in the target phenomenon rather than the other way around. Compare Gillies (2000: 77–84) and Walley (1991: 460–7).

Hume and grue again

We will approach the question of what De Finetti-style subjectivist inductive logic (SIL) has to say about Hume's problem of induction anachronistically, via Goodman's new riddle of induction. Let something be *grue* just in case it is green before some future date D or blue after D. Thus grass is grue before D, but not afterwards, and so on. Goodman thinks that we all agree that "Regularity in greenness confirms the prediction of further cases; regularity in grueness does not" (Goodman 1983: 82). Put probabilistically and in terms of singular predictive inferences: observing green individuals leading up to D-Day raises the probability that the next observed individual (i.e. on or after D-day) will be green whereas it does not raise the probability that it will be grue. Goodman wonders about the basis for distinctions of the green/grue kind. After reproaching Hume for failing to provide such a basis, Goodman himself offers an account that stresses the asymmetrical rootedness of the predicate 'green' in past linguistic and inductive practice (Goodman 1983: chapter 4). Let us now see whether SIL can do as well or better.

In Goodmanian D-day cases, the crucial 'next observed individual' is held fixed at D-day while we haplessly pile up observations prior to its fatefully dated occurrence. An alternative is to treat the next observed individual as a kind of moving target: as we pile up additional observations, the next observed individual, like the proverbial 'free beer tomorrow,' skips ahead always to be observed next. Call these the fixed target and moving target conceptions of 'the next observed individual' respectively (Earman 1992: section 4.7).

Jeffreys (1957) showed that inductive skepticism about the character of the next observed individual in the moving target case ('moving target inductive skepticism') is almost impossible to maintain, since

$$\lim_{n\to\infty} P(Fa_{n+1} \mid Fa_1 \cap \ldots \cap Fa_n) = 1$$

if $P((\forall i)Fa_i) > 0$. This sufficient condition surely has broad skeptical appeal – a skeptic should want to avoid having to be *a priori* certain that $(\exists i)\neg Fa_i$. This is evidently a kind of limiting answer to Hume but it is also a partial answer to Goodman. Since Jeffreys's result does not turn on what 'F' means, contrary to what Goodman might be taken to suggest, grue is on the same footing as green in the moving target sense. No contradiction results from raising the probabilities of both 'the next observed individual is green' and 'the next observed individual is grue' in the moving target sense, since their rates of convergence to the limits in question can, and indeed must be, different (Howson 1973).

Blunting inductive skepticism about the character of the next observed individual in the *fixed* target case ('fixed target inductive skepticism' – clearly the principal case for both Goodman and Hume) requires De Finetti-style symmetry assumptions, not just ringing the changes on the probability calculus. The limit we wish to evaluate in this case (where indices of the a_i now range over both positive and negative integers) is:

$$\lim_{j\to\infty} P(Fa_{n+1} \mid Fa_n \cap \ldots \cap Fa_{n-j})$$

It matters here what 'F' means since suppose day_{n+1} is the D-day for the green/grue divergence and that 'F' means 'is green.' Then

$$\lim_{j\to\infty} P(Fa_{n+1} \mid Fa_n \cap \ldots \cap Fa_{n-j}) = 1$$

implies that

$$\lim_{j\to\infty} P(F^*a_{n+1} \mid F^*a_n \cap \ldots \cap F^*a_{n-j}) = 0$$

where 'F^*' means 'is grue.'

But if P is exchangeable with respect to a given property (i.e. for the infinite sequence of r.v.s comprising the indicator functions for the presence of that property) then the moving target and fixed target limits have to agree. The grue/green case therefore makes for the following inconsistent triad: (1) P is exchangeable with respect to both F and F^*; (2) $P((\forall i)Fa_i) > 0$; and (3) $P((\forall i)F^*a_i) > 0$.

Basic openminded-ness militates against denying either (2) or (3), so exchangeability with respect to at least one of F and F^* must go. Thus, once we grant the inductive skeptic a fixed target at which to aim, symmetries in our subjective probabilities are going to constitute most of our (broadly subjectivist) answer to that skeptic. Those symmetries constitute the respects of resemblance or uniformity that we are expecting to continue into the future, and those determinate expectations implicitly involve us in ignoring countless other abstractly possible respects of resemblance. This is De Finetti's answer to Hume. It is conditional or coherence-minded in much the same way that Goodman's 'past practice' answer is. There's nothing in the theory of exchangeability to say which, if any, properties we should find exchangeable, just as Goodman does not presume to say what our past practices should be. De Finetti's advantage over Goodman is just the clarity afforded within a probabilistic framework for stating and relating the conditions of induction precisely: Bayesian projectibility (Skyrms 1994) is alive and well whereas Goodman's theory of projectibility, as opposed to Goodman's sensational riddle, is a philosophical and logical back-water.

3 New-Fangled Inductive Logic (NFIL)

Logical accounts of the truth-conduciveness of methods of inquiry reached technical and philosophical maturity in the 1980s and 1990s, building on the seminal work of Putnam (1965) and Gold (1965). This New-Fangled Inductive Logic (NFIL) takes guarantees of different senses of convergence to the truth to be the primary object of logical study. Considerations of coherence or consistency – probabilistic or otherwise – are distinctly secondary: they warrant study principally for whether they are likely to block, slow down, or otherwise interfere with convergence to the truth. Our bare-bones treatment of NFIL follows Kelly (1996: chapters 3 and 4).

Consider an idealized scientist trying to determine by passive observation whether some hypothesis, h, is true. We represent the scientist's background knowledge as a set of possible worlds, K, in some of which h is true and in some of which h is false. We suppose that any world in K produces a stream of data, ε, of which the scientist scans only the initial segment, $\varepsilon|n$, up to the current stage, n. The scientist is, we will assume, equipped with an inductive method, α, drawn from some larger set of methods, M, and that the scientist conjectures something about the status of the h after each new data point. We further assume that all of the worlds in K are exhaustively observable. This allows us to identify worlds with their unique data streams, and hypotheses with sets of data streams. Given these identifications, the truth of a hypothesis depends just on the data stream: h is true on ε just in case $\varepsilon \in h$. Lastly, we will assume that the data types are natural numbers and sundry other symbols, which we can think of as codes for more realistic sorts of discrete data types.

Let us now formulate four, increasingly weak senses in which inductive method α may converge to a verdict of some kind.

(C1)	α produces b *by stage n* on h, ε	iff $\alpha(h, \varepsilon\|n) = b$
(C2)	α produces b *with certainty* on h, ε	iff there is a stage n s.t. $\alpha(h, \varepsilon\|n) = !$, $\alpha(h, \varepsilon\|n+1) = b$, and, for all $m < n$, $\alpha(h, \varepsilon\|m) \neq !$
(C3)	α produces b i*n the limit* on h, ε	iff there is a stage n s.t., for all $m \geq n$, $\alpha(h, \varepsilon\|m) = b$
(C4)	α *approaches* b on h, ε	iff for each rational $s \in (0, 1]$, there is a stage n s.t. $\|b - \alpha(h, \varepsilon\|m)\| < s$, for all $m \geq n$

Intuitively speaking, (C1) is convergence by a deadline n, (C2) is convergence to when one is first prepared to say ('!') that one has the answer, (C3) is convergence in the sense of eventual stability in ones conjectures, and (C4) is convergence in the sense of getting closer and closer to some value.

The three most general notions of success for a method α on h, ε involve α conjecturing that h is true ('1') just in case h is true (verification), conjecturing that h is false ('0') just in case h is false (refutation), and both (decision). We provide clauses for the four notions of verification our four senses of convergence induce; clauses for refutation and decision are similar.

(C1v)	α verifies h by stage n on ε	iff	[α produces 1 at n on h, $\varepsilon \leftrightarrow \varepsilon \in$ h]
(C2v)	α verifies h with certainty on ε	iff	[α produces 1 with certainty on h, $\varepsilon \leftrightarrow \varepsilon \in$ h]
(C3v)	α verifies h in the limit on ε	iff	[α produces 1 in the limit on h, $\varepsilon \leftrightarrow \varepsilon \in$ h]
(C4v)	α verifies h gradually on ε	iff	[α approaches 1 on h, $\varepsilon \leftrightarrow \varepsilon \in$ h]

The reliability of an inductive method α is a matter of quantifying over the possible worlds/data streams on which α succeeds, for example:

(C1vK) α verifies h by stage n given K iff for each $\varepsilon \in$ K, α verifies h at n on ε

and so on for the rest. We can also quantify over the range of hypotheses that method α can assess reliably, for example:

(C1vKH) α verifies H at stage n given K iff for each h $\in$ H, α verifies h at stage n given K

Finally, one can ascend to the level of inductive problem solvability by generalizing over the collections of methods in M, for example:

(C3rKHM) H is refutable in the limit given K by a method in M iff there is an $\alpha \in$ M s.t. α refutes H in the limit given K

One can now set about exploring all these notions using the technical palette of the theory of computability and recursion theory. Elementary but important results given no restrictions on M include:

- Verifiability, refutability, and decidability are equivalent for (C1) (Kelly 1996: 45), but not for any of the weaker senses of convergence we have defined (Kelly 1996: 68). For example, the existential hypothesis that mass m is divisible is verifiable with certainty but not refutable (hence not decidable) with certainty, and the same existential within the scope of a universal (e.g. the hypothesis that mass m is infinitely divisible) is refutable in the limit but not verifiable (hence not decidable) in the limit.
- The whole structure can be characterized topologically, roughly by mapping existential and universal hypotheses onto open and closed sets respectively (Kelly 1996: 85; see also Schulte and Juhl 1996).

- Decidability in the limit and gradual decidability are equivalent (Kelly 1996: 67). The class of hypotheses that are so decidable can be classified in the finite Borel hierarchy as Δ_2^B.

NFIL affords an important, abstract yet flexible perspective on inductive inference. It reinvigorates the question of whether conditional, coherence-based answers to Humean and other skepticisms, are answers at all. If the coherencies insisted upon block us from reliably getting to the truth there's a clear sense in which they aren't. NFIL also provides some external check on what might otherwise look like relatively innocuous or 'merely technical' assumptions. Consistency, countable additivity, consideration just of properties (monadic predicates), and so on, are often adopted fairly peremptorily both inside and outside of the GOFIL/SIL tradition. But such assumptions may have hidden costs or be providing illicit benefits, and NFIL promises to help us see clearly whether this is so (Earman 1992: chapter 9; Kelly 1996: chapter 13). NFIL also has the peculiarly philosophical virtue of making strange bedfellows: from its perspective SIL approaches tend to look much more of a piece with traditional, justification-centered programs in epistemology than is usually allowed (see Earman 1992: 219). Our own view is that the articulation of the NFIL perspective on induction – a perspective for which Ramsey calls in the final two sections of his "Truth and Probability" – is a very positive development, but only time will tell whether this view is correct. In any case, both in its ingenuity and its contentiousness, we expect the future of inductive logic to resemble its past.

References

Carnap, R. (1952) *The Continuum of Inductive Methods*. Chicago: University of Chicago Press.

Carnap, R. (1962) *The Logical Foundations of Probability*, 2nd edn. Chicago: University of Chicago Press.

Carnap, R. (1963) The aim of inductive logic. In E. Nagel, P. Suppes and A. Tarski (eds.), *Logic Methodology, and Philosophy of Science: Proceedings of the 1960 International Congress* (pp. 303–18). Stanford, CA: Stanford University Press.

Carnap, R. (1964) Replies and systematic expositions. In P. A. Schillp (ed.), *The Philosophy of Rudolph Carnap* (pp. 859–1013). La Salle, IL: Open Court.

Carnap, R. (1971) A Basic system of inductive logic, part I. In Carnap and Jeffrey (eds.), *Studies in Inductive Logic and Probability* (pp. 33–165). Berkeley, CA: University of California Press.

Carnap, R. (1980) A Basic system of inductive logic, part II. In Jeffrey (ed.), 1980 (pp. 7–155).

Constantini, D. (1983) Analogy by similarity. *Erkenntnis*, 20, 103–14.

De Finetti, B. (1937) Foresight: its logical laws, its subjective sources. In H. E. Kyburg and H. E. Smokler (eds.), *Studies in Subjective Probability* (pp. 93–158). New York: John Wiley & Sons. (Originally published as *La Prévision: ses lois logiques, ses sources subjectives. Annales de l'Institut Henri Poincaré*, 7, 1–68.)

De Finetti, B. (1938) On the condition of partial exchangeability. In Jeffrey (ed.), 1980 (pp. 193–205). (Originally published as *Sur la condition d'équivalence partielle. Actualités scientifiques et industrielles*, 739. Paris: Herman & Cie.)

Diaconis, P. (1977) Finite forms of De Finetti's theorem on exchangeability. *Synthese*, 36, 271–81.

Diaconis, P. and Freedman, D. (1980) De Finetti's theorem for Markov chains. *Annals of Probability*, 8, 115–30.

Diaconis, P. and Freedman, D. (1986) On the consistency of Bayes estimates. *Annals of Statistics*, 14, 1–26.

Dubins, L. and Freedman, D. (1979) Exchangeable processes need not be mixtures of independent identically distributed random variables. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebeite*, 48, 115–32.

Festa, R. (1993) *Optimum Inductive Methods*. Dordrecht: Kluwer.

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

Gillies, D. (2000) *Philosophical Theories of Probability*. London: Routledge.

Gold, E. M. (1965) Limiting recursion. *Journal of Symbolic Logic*, 30, 27–48.

Hájek, A. (1997) 'Mises redux' – redux: fifteen arguments against finite frequentism. *Erkenntnis*, 45, 209–27.

Hewitt, E. and Savage, L. (1955) Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society*, 80, 470–501.

Hintikka, J. (1966) A two-dimensional continuum of inductive methods. In J. Hintikka and P. Suppes (eds.), *Aspects of Inductive Logic* (pp. 113–32). Amsterdam: North-Holland.

Hintikka, J. and Niiniluoto, I. (1980) An axiomatic foundation for the logic of inductive generalizations. In Jeffrey (ed.), 1980 (pp. 157–81).

Howson C. (1973) Must the logical probability of laws be zero? *British Journal of the Philosophy of Science*, 24, 153–63.

Jeffrey, R. (1980) Introduction. In Jeffrey (ed.), 1980 (pp. 1–6).

Jeffrey, R. (ed.) (1980) *Studies in Inductive Logic and Probability*, vol. 2. Berkeley, CA: University of California Press.

Jeffreys, H (1957) *Scientific Inference*, 2nd edn. Cambridge: Cambridge University Press.

Jeffreys, H. and Wrinch, D. (1919) On certain aspects of the theory of probability, *Philosophical Magazine*, 38, 715–31.

Keynes, J. M. (1921) *A Treatise on Probability*. London: Macmillan.

Kuipers, T. (1984) Two types of analogy by similarity. *Erkenntnis*, 21, 63–87.

Maher, P. (2000) Probabilities for two properties. *Erkenntnis*, 52, 63–91.

Niiniluoto, I. (1981) Analogy and inductive logic. *Erkenntnis*, 16, 1–34.

Putnam, H. (1965) Trial and error predicates and a solution to a problem of Mostowski. *Journal of Symbolic Logic*, 30, 49–57.

Salmon, W. (1967) *The Foundations of Scientific Inference*. Pittsburgh: University of Pittsburgh Press.

Schulte, O. and Juhl, C. (1996) Topology as epistemology. *Monist*, 79, 141–7.

Skyrms, B (1993) Analogy by similarity in hyperCarnapian inductive logic. In J. Earman et al. (eds.), *Philosophical Problems of the Internal and External Worlds* (pp. 273–82). Pittsburgh: University of Pittsburgh Press.

Skyrms, B. (1994) Bayesian projectability. In D. Stalker (ed.), 1994 (pp. 241–62).

Tanner, M. A. (1996) *Tools for Statistical Inference: methods for exploration of posterior distributions and likelihood functions*, 3rd edn. Heidelberg: Springer.

Walley, P. (1991) *Statistical Reasoning with Imprecise Probabilities*. London: Chapman & Hall.

Zabell, S. (1992) Predicting the unpredictable. *Synthese*, 90, 205–32.

Zabell, S. (1997b) Confirming universal generalizations. *Erkenntnis*, 45, 267–83.

Further Reading

Earman, J. (1992) *Bayes or Bust: a critical examination of confirmation theory*. Cambridge, MA: MIT Press.

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

Howson C. and Urbach P. (1993) *Scientific Reasoning: The Bayesian Approach*, 2nd edn. La Salle, IL: Open Court.

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

Kuipers, T. (1978) *Studies in Inductive Probability and Rational Expectation*. Dordrecht: Reidel.

Jeffrey, R. (1992) *Probability and the Art of Judgment*. Cambridge: Cambridge University Press.

Maher, P. (1999) Inductive logic and the ravens paradox. *Philosophy of Science*, 66, 50–70.

Ramsey, F. P. (1926) Truth and probability. In D. H. Mellor (ed.), *F. P. Ramsey: Philosophical Papers* (pp. 52–94). Cambridge: Cambridge University Press.

Skyrms, B. (1984) *Pragmatics and Empiricism*. New Haven, CT: Yale University Press.

Stalker, D. (1994) *Grue!* La Salle, IL: Open Court.

Zabell, S. (1988) Symmetry and its discontents. In W. Harper and B. Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics*, vol. II (pp. 155–90). Dordrecht: Kluwer.

Zabell, S. (1989) The rule of succession. *Erkenntnis*, 31, 283–321.

Zabell, S. (1997a) The continuum of inductive methods revisited. In J. Earman and J. Norton (eds.), *The Cosmos of Science* (pp. 351–85). Pittsburgh: University of Pittsburgh Press.