YAEL COHEN

RAVENS AND RELEVANCE

I propose here a new solution to the Paradox of the Ravens, which draws from the two main schools of thought on the paradox – the probabilistic and the Popperian – but avoids their drawbacks. The exposition of this solution makes use of a new presuppositional language which will be outlined in the sequel.

PRESENTATION OF THE PARADOX

The paradox arises out of two principles. The first (called Nicod's Rule) is that all evidence of the form "*Ra* & *Ba*" confirms hypotheses of the form "All *R*s are *B*s". The second principle (called the Equivalence Rule) is that any evidence confirming a given hypothesis also confirms any logically equivalent hypothesis. Thus,

*H*1: "All ravens are black"

and

*H*2: "All nonblack things are nonravens"

are considered logically equivalent hypotheses and, therefore, by these two principles, not only a black raven but also a thing neither black nor a raven confirms *H*1. For example, a white boot (neither black nor a raven) is confirming evidence for "All ravens are black", a highly counterintuitive conclusion.

The logical equivalence of hypotheses *H*1 and *H*2 follows from their representation in standard logic by $(x)(x$ is not a raven $\vee x$ is black). The paradoxical feature arising out of this representation can be highlighted by the following hypothetical situation. A researcher places an advertisement calling for evidence for or against this formal hypothesis. He receives three kinds of evidence: (1) ***a*** is a raven and ***a*** is black, (2) ***b*** is nonblack and ***b*** is a nonraven, and (3) ***c*** is black and ***c*** is a nonraven. Each of these types of evidence verifies a substitution instance of the (formal) hypothesis. The paradox is that there is no way formally to distinguish among the different types of evidence. The

Erkenntnis **26** (1987) 153–179.

differences among them, it is claimed, lie outside logic and depend upon our general knowledge.

THE LOGICO-INDUCTIVE (PROBABILISTIC) SOLUTION[1]

The solution to the paradox furnished by the theory of subjective probability grows out of the quantitative concept of confirmation first developed by Carnap (1950), though exploited earlier in a solution to the paradox by Lindenbaum (1940). The approach was adopted by Gaifman (1979) who offered a fully developed and mathematically rigorous solution. Omitting the mathematics, the solution is that the central question is not *whether* but *how much* an evidence report confirms an hypothesis. Probabilistic calculations show that, despite the existence in theory of three types of evidence to confirm the ornithologist's hypothesis, in practice only type (1) is of any weight, the number of nonblack objects being immeasurably greater than the number of ravens.

It can be shown that probabilistic calculations are unable to account for the confirmatory role of type (3) evidence, and they even indicate that it decreases the degree of confirmation of the hypothesis. But in any event, according to the probabilistic analysis, the contribution of a single nonblack nonraven (which constitutes the basis for type (2) evidence) is miniscule but not zero. The question naturally arises, then, why should a great quantity of evidence concerning objects that are neither ravens nor black not help confirm the hypothesis significantly? Why don't ornithologists, then, include in their observation reports evidence of type (2)? Is it just that the number of nonblack objects is infinite, so that evidence of type (2) contributes no support to the hypothesis? It seems not. Suppose the domain of discourse includes 10^{10} objects of which it is known that exactly ten of them are ravens, and that all 10^7 objects examined were nonblack and nonravens. Even then we would deny that the examined objects support the hypothesis "All ravens are black".

The probabilistic solution suffers from two other drawbacks: it relies on a problematic concept of 'the a priori probability of an hypothesis' and it pictures the scientist as someone who draws objects randomly from the world or from a random sample thereof. According to the probabilistic solution, random sampling is part and parcel of confirmation by instances of strict generalizations.[2]

A POPPERIAN SOLUTION[1]

A different solution to Hempel's 'Paradox of the Ravens', rooted in Popper's conception of evidential support, can be found in Watkins (1958, 1960). According to this conception, boots, for example, are not 'potential falsifiers' of the hypothesis "All ravens are black" and, therefore, cannot confirm the hypothesis. Any *test* of the hypothesis must not preclude its refutation. But why cannot white boots be potential falsifiers relative to a research program according to which we look for nonblack objects and check if they are ravens? This policy could generate a counterexample (if any) to the hypothesis and, therefore, passes Popperian muster. True, the policy is inefficient (given the large number of nonblack objects relative to the number of ravens), but nevertheless it is likely to unearth a counterexample of the hypothesis. Perhaps it is hard to believe that a researcher could have suspected a white object which later turns out to be a boot of being a raven. But this psychological claim is irrelevant to the logic of justification, as distinct from the logic of discovery.

A PRESUPPOSITIONAL LANGUAGE

In this section I present (with no discussion[3]) a presuppositional language – a language extending classical first-order language on whose sentences a relation of presupposition can be defined. The semantics of this language is three-valued, so that it is possible not to assign a truth-value to some sentences. What makes this language important for our discussion is that it contains a presuppositional operator which can also be regarded as a relevance operator, and relevance is a key concept in what follows.

Let L be a first-order language, containing in addition to the usual logical operators ($-$, $\vee$, $\&$, $\rightarrow$, $\leftrightarrow$, ( ), E) two monadic sentential connectives: $'$ and $*$.

A model for L is a quadruple $\langle L, D, S, V\rangle$, where D is a nonempty domain of objects over which the variables range; S is an interpretation function, associating each n-place primitive predicate with an n-place relation on D and each individual term with an object in D; V is a valuation function that assigns to the *wffs* of L values T (true), F (false), or I (undefined) according to the following rules:

For any atomic formula $\psi(v_1, \ldots, v_n)$, V assigns T if and only if $\langle S(v_1), \ldots, (S(v_n)\rangle \in S(\psi)$; otherwise it assigns F.

For compound formulas V assigns values T, F, or I according to the following truth tables: For the standard connectives these are Kleene's 'weak' rules:

p	$-p$
T	F
I	I
F	T

	$p \& q$			$p \vee q$			$p \rightarrow q$			$p \leftrightarrow q$		
$p \backslash q$	T	I	F	T	I	F	T	I	F	T	I	F
T	T	I	F	T	I	T	T	I	F	T	I	F
I	I	I	I	I	I	I	I	I	I	I	I	I
F	F	I	F	T	I	F	T	I	T	F	I	T

That is, V assigns the value I iff at least one of the constituent formulas is assigned the value I; otherwise it behaves standardly. And for the new monadic connectives:

p	$'p$
T	F
I	T
F	T

p	$*p$
T	T
I	I
F	I

The first new connective is interpreted as external negation and is not going to play an essential role in what follows; the second, as we shall see, is interpreted both as a presuppositional operator and as a relevance operator.[4]

For universal formulas V is defined by the following rule:

$V((x)\psi x) = T$ if there is an interpretation function S',

different from S at most with regard to x, according to which $V(\psi x) = T$ and if there is no such S'' according to which $V(\psi x) = F$;

$V((x)\psi x) = F$ if there is an interpretation function S',

different from S at most with regard to x, according to which

$V(\psi x) = F$; otherwise $V((x)\psi x) = I$.

$V(Ex\psi x)$ is defined as $V(-(x) - \psi x)$.

Thus the interpretation of the quantifiers is not standard, e.g., a universal sentence is true iff it has no false instance and has at least one true instance (i.e., it might have instances with undefined truth-value).

According to this definition of a model for L, *the relation of semantic consequence* can be defined as follows:

$A \Vdash B$ if in every model in which A is true, B is also true.

Likewise, *the relation of presupposition* may be defined as:

B is a presupposition of A if $A \Vdash B$ and $-A \Vdash B$.

THE PRESUPPOSITIONAL OPERATOR '*'

Strawson argues that "All John's children are asleep" presupposes that John has children, while "The King of France is bald" presupposes that France has a king. If these presuppositions are false the corresponding sentences are truth-valueless. Strawson generalizes as follows:

> There are many ordinary sentences beginning with such phrases as 'All . . .', 'All the . . .', 'No . . .', 'None of the . . .', 'Some . . .', 'Some of the . . .', 'At least one . . .', 'At least one of the . . .' which exhibit, in their standard employment, parallel characteristics to those I have [. . .] described in the case of a representative 'All . . .' sentence. That is to say, the existence of members of the subject-class is to be regarded as presupposed [. . .] by statements made by the use of these sentences; to be regarded as a necessary condition, not of the truth simply, but of the truth or falsity, of such statements. (1952, p. 176)

Accordingly, in order to formally present the presupposition relation what we need is a presupposition operator, to be prefixed to the predicate denoting the subject-class, causing the whole presupposition-carrying sentence to be true-or-false only if that predicate is true of at least one object. The *-operator as interpreted by the proposed semantics fulfills this requirement exactly. To become convinced of this, consider the following easily verified theorems involving the *-operator:

$$\mathrm{E}x({*}Rx \,\&\, Bx) \vee -\mathrm{E}x({*}Rx \,\&\, Bx) \Vdash \mathrm{E}xRx$$

$$(x)({*}Rx \rightarrow Bx) \vee -(x)({*}Rx \rightarrow Bx) \Vdash \mathrm{E}xRx$$

$$\mathrm{E}xRx \Vdash \mathrm{E}x({*}Rx \,\&\, Bx) \vee -\mathrm{E}x({*}Rx \,\&\, Bx)$$

$$\mathrm{E}xRx \Vdash (x)({*}Rx \rightarrow Bx) \vee -(x)({*}Rx \rightarrow Bx)$$

$$\mathsf{E}x(*Rx \,\&\, Bx) \Vdash \mathsf{E}xRx,\ \mathsf{E}x(Rx \,\&\, Bx)$$

$$-\mathsf{E}x(*Rx \,\&\, Bx) \Vdash \mathsf{E}xRx,\ -\mathsf{E}x(Rx \,\&\, Bx)$$

$$(x)(*Rx \rightarrow Bx) \Vdash \mathsf{E}xRx,\ (x)(Rx \rightarrow Bx)$$

$$-(x)(*Rx \rightarrow Bx) \Vdash \mathsf{E}xRx,\ -(x)(Rx \rightarrow Bx).$$

DIFFERENT HYPOTHESES DESCRIBING THE SAME REQULARITY

Hempel's Paradox of the Râvens is an argument with two premises and an apparently paradoxical conclusion. The present solution modifies both premises and blunts the point of the conclusion.

The second premise of the argument is the Equivalence Rule: what supports a given hypothesis supports any of its logical equivalents. The rule reflects our intuition that support for an hypothesis increases our confidence that the regularity described by it holds, and since logically equivalent hypotheses describe the same regularity, their supporting evidence should be the same.

However, in my system, $H1$: "All ravens are black" and $H2$: "All nonblack things are nonravens" are not *strictly equivalent* (have the same value T, F, I in every model), due to their different existential presuppositions. Moreover, although $H1$ and $H2$ describe the same regularity – the nonexistence of nonblack ravens – they have different logical forms. The two hypotheses have different senses and incorporate different procedures for testing the regularity they describe.

Taking into account the existential presuppositions associated with general statements, $H1$ is represented as:

$$(x)(*Rx \rightarrow Bx)$$

(or the logically equivalent $(x)(*Rx \,\&\, Bx)$, as the reader may prove by himself) and $H2$ as:

$$(x)(*-Bx \rightarrow -Rx)$$

(or the logically equivalent $(x)(*-Bx \,\&\, -Rx)$) where 'Rx' corresponds to 'x is a raven' and 'Bx' to 'x is black'.

Although other representations seem to be permissible, I shall now argue that the recommended representations should be given priority in virtue of how we intuitively grasp the procedures for verifying, refuting and supporting these hypotheses.

Intuitively, it seems that the evidence relevant to hypothesis $H1$ concerns ravens only, while that relevant to $H2$ concerns nonblack things. This intuition seems to warrant Nicod's Rule, on the one hand, and to discredit the conclusion of the paradox, on the other. If we represent these hypotheses as I recommend, we shall be able to articulate the sound element of this intuition by the following rule: any supporting (refuting) evidence gathered via the testing procedure incorporated in an hypothesis like $H1$ or $H2$ ensures the truth (falsity) of at least one substitution instance of the representing sentence. Since the substitution instances of $H1$ having truth-value (T or F) differ from those of $H2$, the supporting evidence gathered via their corresponding procedures will also be different. Every substitution instance of $H1$ is of the form: $(*Ra \rightarrow Ba)$, which has a truth-value if and only if a is a raven. Every substitution instance of $H2$ is of the form: $(*-Ba \rightarrow -Ra)$, which has a truth-value if and only if a is nonblack. In other words, the supporting evidence gathered via the procedure incorporated in $H1$ concerns ravens, while that gathered via the procedure of $H2$ concerns nonblack objects. Nevertheless, in accordance with the rule that a general hypothesis is refuted when at least one of its substitution instances is false, the hypotheses $H1$ and $H2$ are refuted by the same evidence, namely, by nonblack ravens. Any refuting evidence gathered via the procedure incorporated in one hypothesis could be gathered via the procedure of the other.

Our presuppositional language thus allows us to distinguish between $H1$ and $H2$. The two hypotheses are not logically equivalent in our system, because there is a possible state of affairs in which one lacks truth-value while the other is true. In other words, $H1$ presupposes the existence of ravens, while $H2$ presupposes the existence of nonblack objects. Nevertheless, where both presuppositions hold, that is, where there are both ravens and nonblack things, both hypotheses will be true (false) together. But even then the logical form of each hypothesis distinguishes it with respect to its recommended type of supporting evidence: the possibly true substitution instances of each hypothesis relate to different types of objects. The fact that the two hypotheses incorporate different kinds of testing procedures is expressed in the formal language by prefixing the operator '*' to a different predicate. The presuppositional operator thus serves as a relevance operator as well. It is prefixed to the predicate 'x is a raven' in $H1$ because the objects relevant to the testing procedure incorporated in "All ravens are black" include only ravens; it is prefixed to the predicate 'x is nonblack', in $H2$, because the objects relevant to the testing procedure incorporated in "All nonblack things are non-ravens" include only nonblack things.

Hence, though it is true that both hypotheses describe the same regularity (i.e., the nonexistence of nonblack ravens), they point to different ways of establishing this regularity. Using Fregean terms: whenever their presuppositions hold, the two hypotheses have the same referent (truth-value), but different senses; that is, they express two different ways to determine that truth-value.

DEFINITION OF CONCEPTS

The key concepts of the preceding section can be precisely defined in terms of the proposed presuppositional language. It will be assumed that the only names the language contains are canonic names for each object in the world. Our discussion will be limited to contingent hypotheses of the form '$(x)\psi x$', where 'ψx' is an open unquantified sentence. An hypothesis is contingent *iff* there is at least one model in which it is true, and at least one in which it is false.

Hypothesis H' describes the same regularity as hypothesis H'', if:

$$-H' \Vdash -H'' \text{ and } -H'' \Vdash -H'.$$

Thus, $H1$ and $H2$ describe the same regularity.

The domain of relevance of hypothesis H is the set $D = \{x \mid Dx\}$, where 'Dx' is an open sentence satisfying, for every name 'a':

$$Da \vdash (\psi a \vee -\psi a) \text{ and } (\psi a \vee -\psi a) \vdash Da.$$

In other words, the domain of relevance of H is the collection of all the objects in the world for which ψx is either true or false. Thus, the domain of relevance of $H1$ is all ravens and that of $H2$, all nonblack things.

The testing procedure incorporated in hypothesis H is: "Examine the objects belonging to the relevance domain of H (and determine for any examined object whether it satisfies ψx or falsifies it)".

ADEQUACY CONDITIONS FOR A SOLUTION

In light of the above discussion the following revisions of the Equivalence Rule and Nicod's Rule are called for:

The Revised Nicod's Rule: Following the testing procedure incorporated in an hypothesis of the form "All *R*s and *B*s" any supporting evidence will be of the form "*Ra* & *Ba*".

The Revised Equivalence Rule: hypotheses describing the same regularity are supported (falsified) by the same evidence.

The first revised rule does not reflect one of the alleged intuitions embedded in the original rule: it does not suggest that *every* report of the form "*Ra* & *Ba*" constitutes supporting evidence for "All *R*s are *B*s". Yet, these rules still leave us with a seemingly counterintuitive conclusion: a nonblack nonraven (e.g., a white boot) *might* support "All ravens are black". An adequate solution to the paradox should provide us, then, with a notion of *evidence supporting an hypothesis* that would account for the two revised rules, show the discarded intuition to be unwarranted and show that the above conclusion is not untenable. Moreover, a complete solution to the paradox should indicate what testing procedure, among those incorporated in hypotheses describing a given purported regularity, is to be preferred for establishing that regularity. Why, for example, we prefer to establish the regularity stated by both *H*1 and *H*2 through the procedure incorporated in *H*1 and not through that incorporated in *H*2.

MY PROPOSED SOLUTION–INTRODUCTORY DISCUSSION

Let us restrict our attention to a finite world, so that we will be able to speak freely about the relative sizes of different sets of objects. The legitimacy of this restriction will be discussed later.

There are two intuitive requirements on supporting evidence gathered via a procedure incorporated in a given hypothesis. The first, already mentioned, is that ***supporting evidence must verify at least one substitution instance of the hypothesis***, whose form is to be determined according to the proposed formal representation. The second requirement is that ***the estimated number of 'potential falsifiers' of the hypothesis must decrease due to the evidence***. At first glance, this second requirement would seem to follow from the first. However, as we shall see, the estimated number of potential falsifiers of an hypothesis is dependent upon empirical data that may indicate that the estimated number of unexamined potential falsifiers of the hypothesis is smaller than the estimated number of its substitution instances having truth-value (that is, the estimated number of objects in its relevance domain). The explanation why we prefer the procedure of hypothesis $H1$ over that of $H2$ lies in this possibility.

Before rigorously defining the concept of supporting evidence for an hypothesis, let us consider an example. Suppose that we estimate on basis of information at hand that the number of nonblack objects is $m2$, that the number of ravens is $m1$ and that $m1 < m2$ (likewise, $m1$, $m2 > 0$). Since the number (and character) of actual falsifiers, if any, to be gathered via the procedures of either $H1$ or $H2$ is the same, the number of potential falsifiers of each hypothesis is estimated as no greater than the number of ravens and no greater than the number of nonblack things. However, according to the data, the number of ravens ($m1$) is estimated to be smaller than the number of nonblack things ($m2$) and, therefore, before beginning to apply the procedures, it is already known that the number of falsifiers is not greater than $m1$. That is, even at this stage, the estimated number of potential falsifiers is smaller than the estimated number of objects to be examined according to the procedure incorporated in $H2$. This number is equal to the estimated number of objects to be examined according to the procedure of $H1$. Relative to the data that $m1 < m2$, and in the light of the second requirement, the discovery of the first nonblack nonraven *does not* provide supporting evidence. This evidence does not lower our estimate of the number of potential falsifiers of $H2$ (and $H1$) and, thus, according to the second requirement, it does not support $H2$ (or $H1$). This evidence did, indeed, decrease the number of objects yet to be examined according to the procedure incorporated in $H2$, but not the estimated number of potential falsifiers of $H2$ (and $H1$). Both before and with the collection of this evidence, the number of potential falsifiers of $H2$ (and $H1$) is estimated as $m1$. However, even when a nonblack nonraven is observed at a stage where it cannot support $H2$ (or $H1$), it nevertheless enlarges the corpus of examined nonblack nonravens which, as soon as it becomes large enough and contains more than $m2 - m1$ different objects, will support the hypotheses. Then, and *only then*, the accumulated evidence shows that the number of falsifiers is smaller than estimated before.

Thus, on the present approach, although a nonblack nonraven – given a sufficiently large amount of like evidence – constitutes supporting evidence, one is advised to check the alleged regularity according with the procedure incorporated in $H1$. Such an hypothesis (namely, $H1$) having a relevance domain estimated as equal to the estimated number of its potential falsifiers has a special standing. Any evidence gathered via the procedure incorporated in it renders the number of its potential falsifiers smaller than the number of its potential falsifiers estimated according to prior information. Given the above prior information, any black raven supports the hypotheses by decreasing the number of their potential falsifiers as estimated according to the prior information, that is, the estimated number of objects in the relevance domain of $H1$. The testing procedure incorporated in $H1$ is more economic than the one incorporated in $H2$ and, therefore, preferred.

WHAT IS SUPPORTING EVIDENCE–THE PROPOSED SOLUTION

The concepts by means of which the conditions governing supporting evidence are formulated can be precisely defined in terms of the presuppositional language described above. The following definitions, together with the preceding ones, provide a complete solution to the paradox.

Prior information concerning a set of hypotheses $\{H_i\}_i$ (to be denoted by $I_{\{H_i\}_i}$) is a sentence which:

(i) implies the bivalence of each member of the set $\{H_i\}_i$ (i.e., their presuppositions hold),
(ii) does not imply the falsity or the truth of any member of the set $\{H_i\}_i$ (i.e., it is still reasonable to try to falsify them),
(iii) implies sentences concerning the classes corresponding to Boolean combinations of primitive predicates (for example, sentences reporting the powers of such classes).

Hypothesis H' describes the same regularity as hypothesis H'' relative to prior information $I_{\{H_i\}_i}$, where $H', H'' \in \{H_i\}_i$, if:

$$I_{\{H_i\}_i} \Vdash (H' \leftrightarrow H'').$$

Thus, since $H1$ and $H2$ describe the same regularity, they a fortiori describe the same regularity relative to any prior information that concerns them. However, $H1$ and $(x)(*(Rx \vee Gx) \rightarrow Bx)$, e.g., describe the same regularity only relative to prior information that implies '$-\mathrm{E}xGx$'.

Henceforth, where prior information $I_{\{H_i\}_i}$ is referred to in a context involving an hypothesis H, it will be assumed that H belongs to $\{H_i\}_i$, and the subindex will be omitted. Likewise, all hypotheses are to be of the form $(x)\psi x$.

The domain of falsifiers of hypothesis H is the set $M = \{x \mid Mx\}$ where 'Mx' is an open sentence satisfying for every name 'a':

$$-\psi a \Vdash Ma \text{ and } Ma \Vdash -\psi a .$$

'Mx' is therefore a predicate satisfied by all and only the objects that falsify H. The domain of falsifiers of $H1$ and $H2$ is, of course, the set of nonblack ravens ('Mx' is 'Rx & $-Bx$').

The power of the relevance domain of hypothesis H estimated according to prior information I (to be denoted by $|D_H|_I$) is the estimated least upper bound of the power of D according to I.[5]

The number of potential falsifiers of hypothesis H estimated from prior information I (to be denoted by $|M_H|_I$) is the estimated least upper bound of the power of M according to I.[5]

Thus, given a prior information telling us just that there are $m1$ ravens and $m2$ nonblack objects, the estimated number of potential falsifiers of $H1$ (and $H2$), that is, the estimated power of the set corresponding to 'Rx & $-Bx$', is $min\{m1, m2\}$. This is the best estimate of the power of the intersection set when all we know is the powers of the intersecting sets.

Now, what I suggest is that supporting an hypothesis amounts to decreasing the estimated number of its potential falsifiers. That is, evidence supports an hypothesis when and only when it demonstrates that the number of falsifiers of the hypothesis is smaller than that estimated according to prior information. Such a conception of evidential support renders a complete solution to the paradox. Let us then define:

Given prior information I, the number of potential falsifiers of H according to evidence e (to be denoted $|M_H|_{I \,\&\, e}$) is the estimated least upper bound of the power of M according to I & e.

Given prior information I, evidence e supports an hypothesis H if:

$$|M_H|_{I\,\&\,e} < |M_H|_I.$$

The degree of support could then naturally be defined as the proportion of the number of potential falsifiers reduced by evidence to the number of potential falsifiers according to the given prior information.

Given prior information I, the degree of support of hypothesis H relative to nonrefuting evidence e is the value of the three-place C function defined as:

$$C(H, e, I) = 1 - \frac{|M_H|_{I\,\&\,e}}{|M_H|_I}$$

Now, since every falsifier of an hypothesis H is a falsifier of every hypothesis describing the same regularity (relative to a given prior information), all hypotheses describing the same regularity (relative to a given prior information) have the same number of potential falsifiers, whether it is estimated according to that information or according to that information enriched with some nonrefuting evidence. Thus, given any prior information, any evidence that supports a given hypothesis also supports any hypothesis describing the same regularity. In other words, hypotheses describing the same regularity are supported (refuted) by the same evidence, as required by the Revised Rule of Equivalence.

Moreover, the set of potential falsifiers of H might be estimated, according to a given prior information, as smaller (but never greater) than its relevance domain. In such a case, to follow the procedure incorporated in H would bring us to lower our estimate of the number of H's potential falsifiers, only when the number of the yet unexamined objects in the relevance domain would become smaller than that prior estimate. Only then the evidence gathered via the procedure of H would constitute supporting evidence, for H and for any other hypothesis describing the same regularity. The suggested revision of Nicod's Rule is thus warranted. The original rule says, in the terminology advocated here, that evidence verifying a substitution instance of a (universal) hypothesis supports that hypothesis. I have argued against this rule where the relevance domain of the hypothesis is estimated to be larger than its domain of potential falsifiers according to prior information. In such a case, it is valid only after enough evidence has been gathered to decrease the prior estimated number of potential falsifiers. In other words, a true substitution instance of an hypothesis should not count as supporting evidence unless it reduces the estimated number of potential falsifiers.

Likewise, although, given what we know, a nonblack nonraven does not support "All ravens are black"; coming on the heels of a sufficiently large number of like objects (i.e., a number which is at least greater than the difference between the number of nonblack things and that of ravens), it does support it. Accordingly, the definitions imply that among hypotheses describing the same regularity (relative to a given prior information), other things being equal, the one, whose relevance domain is estimated to be the smallest, points to the best procedure for examining the regularity described.

We have thus disclosed the source of the unintuitiveness of the paradox's conclusion, namely, that a nonblack nonraven might support "All ravens are black". First of all, evidence about nonblack nonravens is never gathered via the testing procedure incorporated in that hypothesis. Second, we all know that ravens are much fewer than nonblack things (or even than nonblack birds), so even if we choose to establish the regularity in question by the procedure incorporated in "All nonblack things are nonravens", such evidence would have to be extremely large in order to support the hypotheses describing the regularity. The conclusion is not, however, untenable: if you are patient and diligent enough, a nonblack nonraven will indeed support the hypothesis.

PRIOR INFORMATION AND THE PREFERRED PROCEDURE

So far I have referred explicitly to information giving a least upper estimate only on the powers of the domains of ravens and of nonblack things in general. However, given prior information according to which, for example, nonblack birds are fewer than ravens, we may prefer to examine nonblack birds following the procedure of the hypothesis.

"All nonblack birds are nonravens"

which describes the same regularity as $H1$ and $H2$ (taking into account minimal semantic knowledge).

There could also be prior information according to which the preferred procedure is neither that of $H1$, nor that of $H2$ (not even a sub-procedure of these). The situation is analogous to the familiar four-card problem (Wason, 1966). Consider a group of four objects a, b, c, and d, about which the following is known: Ra, $-Rb$, Bc, $-Bd$. The regularity to be examined is, again, whether there are R's which are non-B in the group. According to prior information, the alleged regularity is described by:

$$(x)(*Rx \rightarrow Bx),\ (x)(*-Bx \rightarrow -Rx),\ (x)(Rx \rightarrow Bx),$$
$$(x)(*(-x = b \ \&\ -x = c) \rightarrow (Rx \rightarrow Bx)), \text{ etc.;}$$

the estimated powers of the relevance domains are 3, 3, 4, 2, respectively, and the estimated number of potential falsifiers is 2. Now, although the relevance domain of the third hypothesis is the largest, the prior information provides two true instances (by b and c), so there are just two more objects to be examined to see if the regularity holds. The procedure to be preferred is, then: "Check every object in the group except b and c (which were already found to be irrefuting)". This procedure is incorporated in the fourth hypothesis and constitutes a sub-procedure of that incorporated in the third. Notice that in order to verify the regularity through the first two hypotheses we would have to check three objects (some of which might be found irrelevant). Again, the most economic procedure is the one incorporated in the hypothesis with the smallest relevance domain.

Yet, one should remember that the selection of a testing procedure to test for a regularity is also made according to various other pragmatic considerations. For example, if we know that all ravens are either black or gray, and we want to verify the hypothesis "All ravens that have chromosomes of type α are black" we shall presumably prefer the procedure incorporated in "All gray (nonblack) ravens do not have chromosomes of type α", which describes the same regularity according to the given information. This, even if we know that gray ravens are more numerous than ravens having type α chromosomes (or than the number of black ravens). For, obviously, the colour of a bird is easier to ascertain than the type of its chromosomes.

There are also testing procedures incorporated in hypotheses describing the same regularity as $H1$ which for epistemic reasons will never be chosen. Consider, for example, the hypothesis:

$H3$: "Everything that is a raven, if and only if it is nonblack, is black",

which is represented by

$$(x)(*(Rx \leftrightarrow -Bx) \rightarrow Bx).$$

The non-refuting evidence gathered via the procedure incorporated in $H3$ is of the form $-Ra$ & Ba, that is, evidence about black nonravens. But what does it mean to follow that procedure? It seems that in order to establish whether an object belongs to the relevance domain of $H3$, one must establish simultaneously whether or not it refutes the hypothesis. Accordingly, it is difficult to imagine that we would have prior information relative to which the power of the relevance domain of $H3$ would be estimated as not greater than that of its potential falsifiers. Consequently, for reasons concerning the nature of the property determining the relevance domain of a given hypothesis – the way we ascertain whether or not an object has that property – the procedure of $H3$ will never be preferred. The procedure which may yield supporting evidence exclusively of the type $-Ra$ & Ba is, for general epistemic reasons, one we would never choose.[6]

NATURAL KINDS

There are still other more general considerations in determining the preferred procedure. Ravens constitute what is called a 'natural kind' or 'scientific kind', that is, the population of ravens reveals, according to the accepted conceptual scheme, great regularity. This allows us to classify ravens for the purpose of testing the regularity under discussion according to the properties that an accepted scientific theory links to bird colour, such as habitat, age, sex, etc.

This 'scientific' division into types makes it unnecessary for us to examine all ravens to verify the regularity; it is enough to observe several representatives of each type of raven. We may thus relativize the function C that measures the degree of support of an hypothesis to the estimated number of potential-falsifier *types* (e.g., raven types) and to the number of potential-falsifier *types* (e.g., raven types) whose representatives were examined and found to be nonrefuting (e.g., found to be black). There is, however, no analogous way to diminish the nonblack things to be examined according to $H2$. The nonblack population is not only immeasurably larger than the raven population, but fails to display the sorts of regularities that would allow its division into types.

Thus, even if the number of ravens were as large as that of nonblack things, even then we would prefer to test the regularity that there are no nonblack ravens by examining ravens and not nonblack things. For even then the raven types would be fewer than the nonblack things, and therefore our confidence in the regularity would grow with each observed black raven of a type not yet examined, more than it would upon our observing many nonblack nonravens.

TESTING PROCEDURES INVOLVING RANDOM SAMPLING

So far I have discussed in detail only evidence collected via testing procedure of checking objects belonging to some given relevance domain, where the method by which these objects are collected is assumed to be unspecified. Thus, given prior information I which does not imply (true) instances of H, the number of potential falsifiers of H according to evidence e verifying k instances of H and collected in a statistically unprescribed method (e.g., not by random sampling) is:

$$min(|M_H|_I, |D_H|_I - k).$$

It has accordingly been claimed that, since ravens are much fewer than nonblack things, even many nonblack nonravens will not support the regularity in question, unless their number is greater than the (estimated) difference between the number of nonblack things and the number of ravens.

However, consider a piece of evidence verifying k different substitution instances of H which has been collected by sampling from the relevance domain of H. In that case, it might be possible to estimate the number of potential falsifiers as much lower than that estimated according to similar evidence collected without sampling. Intuitively, a very large random sample of nonblack things found to include only nonravens support $H2$, even where their number is smaller than the (estimated) difference between the number of nonblack things and the number of ravens.

Where neither the prior information nor the evidence essentially involve probability terms, the number of potential falsifiers is the number of (actual) falsifiers estimated so as to logically ensure that no further evidence (consistent with the information at hand) would indicate that it is larger than estimated. Yet, where, for example, the evidence is based on a random sample, one may be content with a cautious and reliable, though not absolutely secure, estimate of the least upper bound of the number of falsifiers. In other words, given evidence that k objects sampled from the relevance domain of H were all found to be non-refuting, the least upper bound of the set of H's falsifiers can be given an estimate which is very unlikely to be found to be too low, though it logically can be so. Now, the appropriate question to be asked, when given such evidence, is the following: Given the prior information, what is the least upper bound of the number of falsifiers of H which presents this evidence as not very improbable? An example might help in clarifying the matter.

Suppose that according to prior information I there are 50 nonblack things and 25 ravens. The number of potential falsifiers of $H2$ (or of $H1$) according to I is 25. Suppose further that a random sample of 25 objects drawn from among the nonblack things has been found to include 25 nonravens. Given the prior information I, what is the number of potential falsifiers according to the evidence?

According to I the least upper bound of the number of (actual) falsifiers of $H2$ (or $H1$) is 25. Adding the given evidence to I does not logically entail a different least upper bound, though it does provide a lower estimate of that bound which is highly reliable from a standard statistical point of view. Adopting classical statistical approach, there is a set of questions the answers to which lead to a statistically reliable estimate of the number of falsifiers. The questions are: for every n, $25 > n \geqslant 0$, assuming that there are n (actual) falsifiers of $H2$, what is the probability that a random sample of size 25 taken from the nonblack population includes no falsifiers of $H2$? When these questions are given answers according to standard probabilistic reasoning, one should choose a ***significance level*** determining what is the smallest value of the above probabilities which should be considered as non-negligible, i.e., as insufficiently small for justifying the rejection of the possibility that the upper bound of the number of falsifiers is smaller than the assumed number of falsifiers yielding that probability. The least upper bound of falsifiers is estimated as the smallest number of two numbers: the number of falsifiers estimated according to the prior information and the largest (assumed) number of falsifiers that yields a probability of the given evidence which is not smaller than the chosen significance level.

Let us then consider the example in detail. Assuming that there are n falsifiers, $n < 25$, the probability p_n of the given evidence – no falsifiers in a random sample of size 25 taken from the population of 50 objects – is:

$$p_n = \frac{\binom{50-n}{25}}{\binom{50}{25}}$$

Obviously, p_n increases monotonously with n. Let us now choose a standard significance level of 1%. It is, therefore, sufficient to note that, in this case, the largest value of n, such that $p_n \geqslant 1\%$, is 6. We can then say that a cautious and highly reliable estimate of the number of falsifiers of $H2$ is 6.[7]

The notion of random sampling is thus naturally incorporated into my general proposal of a solution to Hempel's paradox. The applicability of the notion of random sample to actual evidence can indeed be very problematic, but my general proposal does not depend on this notion. Contrary to the subjective probability solution, the proposed solution does not essentially depend on probability notions. Probability comes into play only when the evidence is already given in terms of random sampling, and even then, the problematic notion of 'a priori probability of an hypothesis' has no role.

'ANY' CLAIMS

There is a pair of sentences closely connected with the pair of hypotheses $H1$ and $H2$:

$H'1$: "Any raven you may select will be black"

and

$H'2$: "Any nonblack thing you may select will be a nonraven".

In case that $H'1$ and $H'2$ are not simply understood as reformulations of $H1$ and $H2$, they are used for making an offer (or a suggestion) followed by a claim (which together constitute a challenge).[8] In saying $H'1$ we make a *suggestion*:

Pick up an object from amongst the ravens

and *claim* that:

The chances of then picking a nonblack one are very low.

In saying $H'2$ we make a *suggestion*:

Pick up an object from amongst the nonblacks

and *claim* that:

The chances of then picking a raven are very low.

Under this interpretation $H'1$ and $H'2$ are clearly nonequivalent. The corresponding suggestions are different: $H'1$ relates to the raven population while $H'2$ relates to the nonblack population (much the same as $H1$ is *about* ravens while $H2$ is *about* nonblacks). Consequently, the corresponding claims are nonequivalent, for clearly one could be warranted without the other being so. In a world where there are millions of nonblack things and just two ravens, one black and one nonblack, only (the claim corresponding to) $H'2$ would be borne out but not (the claim corresponding to) $H'1$. In any world in which ravens are much less numerous than nonblack things (the claim corresponding to) $H'2$ would be borne out, no matter how many, if any, black ravens there are.

Let us say, using our notation, that according to a given information I the number of ravens is estimated as $|R|_I$, that of nonblacks as $|-B|_I$ and that of the potential falsifiers of $H1$ (or of $H2$ or of any other hypothesis describing the same regularity according to I) as $|M|_I$. The claim of $H'1$ is borne out if the probability of picking up from amongst the ravens a nonblack one is very low. This probability is estimated, by simple standard probabilistic reasoning, as: $1/2(|M|_I/|R|_I)$. The same goes for $H'2$; the relevant probability is estimated as $1/2(|M|_I/|-B|_I)$. The lower the estimated number of potential falsifiers, the lower these estimated probabilities. The larger the number of objects in the associated relevance domain, the lower the probability. Thus, any evidence which reduces the estimated number $|M|$ would support both $H'1$ and $H'2$ but to a different degree. Likewise, knowing that $|R|$ is much smaller than $|-B|$ constitutes a sufficient ground for claiming $H'2$ but not for claiming $H'1$.

We see, then, that the two central ideas on which my solution to the paradox is based pertain to $H'1$ and $H'2$ as well. The first is that statements of different logical form may have a common domain of falsifiers, so that any evidence that reduces the estimated number of objects belonging to that domain affects their support. The second is that statements may differ in their relevance domains due to their logical form, so that quantitative estimations of these domains are relevant to their support. The first provides us with a quantitative concept of supporting evidence for hypotheses like $H1$. The second bears upon a further question: how large should evidence concerning objects in the relevance domain be in order to constitute supporting evidence. All this is relevant to $H'1$ and $H'2$ through their connection with $H1$ and $H2$, reflected by the fact that $|M|_I$ appears in both associated probabilities. But in this case, the second idea is of further significance: it indicates that the claims made by $H'1$ and $H'2$ are different since they are not necessarily borne out together (not even when the presuppositions of $H1$ and $H2$ hold).[9]

THE PARADOX AND LAWS OF NATURE

The proposed solution of the paradox bears upon support of lawlike statements as far as induction *by enumeration* is involved in their support.

The proposed solution rests upon one crucial claim and one crucial supposition which seem unjustified when considered with regard to lawlike statements. The claim pertains to the logical form of statements of the type: "All Rs are Bs". I have claimed that *all* such statements have a logical form that indicates an existential presupposition (the existence of Rs) and thereby a definite relevance domain. The existence of Rs is presupposed in the sense that it is a necessary condition for the statement being true-or-false. We know, however, that the lawlike statement "All bodies not acted upon by external forces move at a constant velocity" was held true and the statement "All bodies not acted upon by external forces do not move at a constant velocity" was held false, while nobody believed that their subject-class is non-empty. But this is just another indication of the inapplicability of induction *by enumeration* to such lawlike statements (expressing possible fundamental laws). The first statement could not have been supported, and the second could not have been falsified, by examining bodies not acted upon by external forces. In other words, (the revised) Nicod's rule loses its force with regard to such purported laws and the paradox does not arise.

What is then the logical form of fundamental laws? N. Cartwright (1983) argues that fundamental laws are true only of objects in the model of the theory and that the phenomenological laws (hypotheses like $H1$ are ones) are the only laws which might be true of objects in reality. The distinction between the two kinds of laws is clearly not semantic. It mainly pertains, she argues, to their different functions, explanatory vs. descriptive, and to the different ways in which they might be supported. Thus we may analyse *all* statements of the type "All Rs are Bs" as having the same logical form, and yet distinguish between them according to the type of entities over which their variables range.

The supposition I have made, superimposed on the above claim, is that the relevance domains are closed and finite. The supposition seems unjustified even with regard to phenomenological generalizations like $H1$. Such generalizations, unlike fundamental laws, may be supported through induction by enumeration, yet we believe their domain of relevance to be unbounded. This observation, however, could be reconciled with the supposition. The inductive reasoning that might lead us to accept an hypothesis like $H1$ could be seen as consisting of two major steps. The first step is induction by enumeration of the given hypothesis taken as restricted to some time-interval. The supposition, which is essential only to this step, is then justified and the proposed quantitative considerations are, therefore, applicable. The second step is of nonquantitative nature. It is a complex reasoning based on the 'Principle of the Uniformity of Nature', which is interpreted in terms of the accepted fundamental categories of the science to which the hypothesis belongs.

SUMMARY OF THE SOLUTION

The solution combines two intuitions concerning evidence: first, given different domains in which one might test an alleged regularity, the narrowest is the most preferable and, second, support for an hypothesis consists in excluding potential falsifiers. The probabilistic solution relies primarily upon the first intuition, analyzing it by means of a subjective probability function; the Popperian solution, upon the second. The present analysis combines the virtues of both approaches and avoids their drawbacks.[10]

The second intuition is expressed by the requirement that support for an hypothesis increases as, and just as, the estimated number of potential falsifiers is reduced. Of course, background information (such as, the estimation of the powers of the relevance domains and the prior estimation of the number of potential falsifiers) determines whether a given piece of evidence is supporting or not, but this determination does not depend on how the evidence is obtained. In this way we avoid the psychologistic elements of the Popperian approach but acknowledge the influence of background information.

The first intuition is captured by the following rule: among the testing procedures incorporated in hypotheses purporting to describe the same regularity, choose the one incorporated in the hypothesis whose relevance domain is estimated to be the smallest. This instruction ensures – ignoring additional considerations relating to natural kinds – that the number of objects to be examined in order to demonstrate the regularity will be minimal; and where the relevance domain is estimated as equal to the prior estimation of the number of potential falsifiers, any observed object that does not falsify the hypothesis, supports it. Now, subjective probability theorists make a similar recommendation. On their approach, however, even the first observed nonblack nonraven slightly strengthens our confidence in the Hempelian regularity. On my approach, since ravens are much fewer than nonblack things, even many nonblack nonravens, observed as a result of a statistically unprescribed method, will leave us cold, unless their (estimated) number is greater than the (estimated) difference between the number of nonblack things and the number of ravens. Finally, my solution avoids problematic concepts like 'the a priori probability of an hypothesis', which is essential to the probabilistic approach, and, while leaving room for statistical considerations, it does not make random sampling part and parcel of confirmation by instances of strict generalizations.[11]

NOTES

[1] I have no intention to fully present or discuss either the probabilistic solution or the Popperian one. The presentation is just to serve as a background for my own solution –

stressing the intuitions behind these two approaches which I embrace, and their drawbacks which I avoid.

[2] See, for example, the representative experiment suggested in Gaifman (1979, pp. 108–109, 128).

[3] The full presentation of this presuppositional language appears in my doctoral dissertation.

[4] This set of three-valued connectives is functionally complete.

[5] Given Boolean expression $\beta(A_1, \ldots, A_n)$; the powers of the sets of objects satisfying A_i $(1 \leqslant i \leqslant n)$; and the power of the domain of discourse, there is an algorithm for calculating the estimated least upper bound of the set of objects satisfying $\beta(A_1, \ldots, A_n)$. In other words, given a domain of discourse U and subsets $A_1, \ldots, A_n$, together with their corresponding powers $|U|$, $|A_1|, \ldots, |A_n|$, there is a computable function $f(|U|, |A_1|, \ldots, |A_n|)$, providing the best estimate of the least upper bound of the power of the set corresponding to $\beta(A_1, \ldots, A_n)$; that is,

$$f(|U|, |A_1|, \ldots, |A_n|) = \textit{minimax}\ |\beta(A_1, \ldots, A_n)|$$

where $|\beta(A_1, \ldots, A_n)|$ is the power of the set of objects satisfying $\beta(A_1, \ldots, A_n)$ for subsets $A_1, \ldots, A_n$ of the domain U. Consider, for example, a Boolean combination

$$\beta(A, B, C) = (A \,\&\, B) \vee (B \,\&\, C) \vee (C \,\&\, A)$$

Suppose that the power of the discourse domain U is 1000 and that $|A| = |B| = |C| = 10$. Then our algorithm yields 15 as the best estimate for the lower upper bound of the total number of objects satisfying $\beta(A, B, C)$; that is, the number of such objects is $\leqslant 15$ in any world in which the above data hold (and equals 15 in one such world). I am indebted to Professor Y. Stavi, who showed me how to construct such an algorithm by means of a proof method of quantifiers elimination.

[6] There are seven different (logically nonequivalent) hypotheses describing the same regularity as $H1$:

(1) $(x)(*Rx \rightarrow Bx)$

(2) $(x)(*-Bx \rightarrow -Rx)$

(3) $(x)(*(Rx \leftrightarrow -Bx) \rightarrow Bx)$

(4) $(x)(*(-Bx \vee Rx) \rightarrow (Rx \leftrightarrow Bx))$

(5) $(x)(*(Bx \vee Rx) \rightarrow Bx)$

(6) $(x)(*(-Bx \vee -Rx) \rightarrow -Rx)$

(7) $(x)(-Rx \vee Bx)$

What is peculiar to evidence of the form $-Ra$ & Ba is that the only testing procedure that yields such evidence is the one incorporated in $H3$, which we have seen to be epistemically inadmissible. All the other procedures can be followed without yielding evidence of this type. Some of these procedures (those of (1), (2) and (4)) tell us explicitly to avoid examining black nonravens, and, therefore, clearly cannot yield such evidence. The rest are procedures that either apply to *all* black things or to *all* nonravens. In the first case (that of (5) and (7)), following the procedure, when applied to a black thing, does not require any further determination of the object – any black thing satisfies a substitution instance of the corresponding hypothesis. Likewise in the second case (that of (6) and (7)), following the procedure, when applied to a nonraven, does not require any further determination of the object – any nonraven satisfies a substitution instance of the corresponding hypothesis. Thus, none of the admissible procedures for testing the regularity in question would yield (supporting) evidence of the form $-Ra$ & Ba, and in this sense such evidence is always either wholly or partially redundant.

[7] The general case is the following: According to prior information, which does not imply instances of H, the number of potential falsifiers of H is m and the power of its relevance domain is r, where $m \leq r$. A random sample of size k is taken from the relevance domain and found to include no falsifier of H. The number of potential falsifiers, according to the information at hand and relative to a significance level of 1%, is the smallest of two numbers:

$min(m, r-k)$ and the largest n yielding a probability p_n

not smaller than 1%; the values of p_n are given by:

$$p_n = \frac{\binom{r-n}{k}}{\binom{r}{k}}$$

It might be worth noting that not every random sample of nonfalsifiers reduces the number of potential falsifiers. Consider, for example, prior information according to which the relevance domain of H includes 50 objects and the number of potential falsifiers is 25. The number of potential falsifiers, given additional nonrefuting evidence consisting of a random sample of 5 objects taken from the relevance domain, can be shown to be still 25.

[8] The general idea is to be found in Z. Vendler (1962).

[9] Similar reasoning applies to hypotheses like: "Almost every raven is black" and "Almost every nonblack thing is nonraven".

[10] Mackie (1963) also tries to reconcile the probabilistic approach to the paradox with the Popperian one. For a context in which the only background information is that ravens are less numerous than nonravens, and black things than nonblack things, he proposes a simplistic version of the probabilistic solution. His solution suffers from *all* the drawbacks of the probabilistic solution I mentioned. Moreover, Mackie's reason for the claim that a nonblack nonraven confirms $H1$ to a lesser degree than a black raven, is inconclusive. One should explain why the contribution of a given observation report is to be the *ratio* of its probability in relation to the background information enriched by the hypothesis to its probability in relation to the background information alone. Why not, for example, the *difference* between these probabilities, in which case the contributions of both kinds of reports would appear to be the same?

Mackie brings in Popperian ideas when discussing richer background information. The background information is enriched by specifying the policy of the investigator, and, thus, different policies are supposed to affect differently the above ratio of probabilities. According to this reasoning, Mackie maintains, the best policy is: "Look for nonblack ravens", which is the policy recommended by Popperians. However it is very unclear how this preferred policy could differentiate between the two kinds of reports. Moreover, what does it mean to adopt this policy? The only clarification Mackie gives us is that it is a policy which highly raises our chances of falsifying $H1$. But this clarification does not help much. What is a policy which highly raises our chances of falsifying $H1$? This is something an adequate solution to the paradox should not take as immediately understood, but should instead explain in detail.

Mackie considers three policies allegedly available for examining $H1$: "Look for black ravens", "Do not look for nonblack ravens" and "Look for nonblack ravens". Mackie does not explain how are we to go about when adopting one of these policies. In the light of my proposed logical analysis, these policies might be explained as incorporated, respectively, in hypotheses of the form:

$$(x)(*(Rx \mathbin{\&} Bx) \rightarrow \cdots),$$

$$(x)(*(-Rx \vee Bx) \rightarrow \cdots), \text{ and}$$

$$(x)(*(Rx \mathbin{\&} -Bx) \rightarrow \cdots).$$

None of these hypotheses, however, describes the same regularity as $H1$, and, therefore, none of these policies is appropriate for checking the alleged regularity. One may, of course, regard any hypothesis describing the same regularity as $H1$ as incorporating a procedure for finding nonblack ravens, if there are any, but an adequate solution to the paradox should indicate which of these procedures are to be preferred and why.

[11] I would like to thank Prof. M. Steiner, Prof. J. Stern, Dr. G. Bareli, Dr. D. Wiederker and Dr. T. Kulka for their helpful comments. I am also indebted to two anonymous referees of *Erkenntnis* for their very helpful criticism of a previous version of this paper, and especially to Prof. Yonathan Stavi for his invaluable help in writing this paper.

REFERENCES

Carnap, R.: 1950, *Logical Foundation of Probability*, Chicago University Press, Chicago, 2nd ed., 1962.

Cartwright, N.: 1983, *How the Laws of Physics Lie*, Clarendon Press, Oxford.

Gaifman, H.: 1979, 'Subjective Probability, Natural Predicates and Hempel's Ravens', *Erkenntnis* **14**.

Linderbaum-Hasiasson, J.: 1940, 'On Confirmation', *Journal of Symbolic Logic* **5**.

Mackie, J. L.: 1963, 'The Paradox of Confirmation', *The British Journal for the Philosophy of Science* **13**.

Strawson, P. F.: 1952, *Introduction to Logical Theory*, Methuen.

Wason, S.: 1966, 'Reasoning', *New Horizons in Psychology*, 1, (B. Foss (ed.), Penguin). See discussion in Cohen J. L., 1981, 'Can Irrationality Be Experimentally Demonstrated?', *The Behavioral and Brain Sciences* **4**, 317–331.

Watkins, J. W. N.: 1958, 'A Rejoinder to Hempel's Reply', *Philosophy* **33**.

Watkins, J. W. N.: 1960, 'Conformation Without Background Knowledge', *The British Journal for the Philosophy of Science* **10**.
Vendler, Z.: 1962, 'Each and Every, Any and All', *Mind* **71**.

Manuscript submitted 7 August 1985
Final version received 8 April 1986

Faculty of Humanities
Department of Philosophy
The Hebrew University of Jerusalem
Jerusalem
Israel