6.2

The Raven Paradox

Carl Hempel

According to the received view, scientific laws are universal or statistical generalizations of the form "all (or most) As are Bs," such as "all ravens are (creatures that are) black." The sighting of a black raven would intuitively confirm this law. But, Hempel points out in this selection, if one assumes a very intuitive principle – that an observation that confirms a generalization also confirms any other statement that is logically equivalent to that generalization – it appears that a red pencil also confirms the law that all ravens are black, which is highly unintuitive. Hempel's paradox presents a serious obstacle to the attempt to formulate a purely formal theory of confirmation.

1 Nicod's Criterion of Confirmation and Its Shortcomings

We consider first a conception of confirmation which underlies many recent studies of induction and of scientific method. A very explicit statement of this conception has been given by Jean Nicod in the following passage: "Consider the formula or the law: *A entails B*. How can a particular proposition, or more briefly, a fact, affect its probability? If this fact consists of the presence of B in a case of A, it is favorable to the law: '*A entails B*'; on the contrary, if it consists of the absence of B in a case of A, it is unfavorable to this law. It is conceivable that we have here the only two direct modes in which a fact can influence the probability of a law.... Thus, the entire influence of particular truths or facts on the probability of universal propositions or laws would operate by means of these two elementary relations which we shall call *confirmation* and *invalidation*."[1] Note that the applicability of this criterion is restricted to hypotheses of the form '*A entails B*'. Any hypothesis H of this kind may be expressed in the notation of symbolic logic[2] by means of a universal conditional sentence, such as, in the simplest case,

$$(x)[P(x) \supset Q(x)]$$

i.e. 'For any object x: if x is a P, then x is a Q,' or also 'Occurrence of the quality P entails occurrence of the quality Q.' According to the above

From Carl G. Hempel, "Studies in the Logic of Confirmation (I)," *Mind*, Vol. LIV, No. 213 (1945): 1–26.

criterion this hypothesis is confirmed by an object a if a is P and Q; and the hypothesis is disconfirmed by a if a is P, but not Q.[3] In other words, an object confirms a universal conditional hypothesis if and only if it satisfies both the antecedent (here: '$P(x)$') and the consequent (here: '$Q(x)$') of the conditional; it disconfirms the hypothesis if and only if it satisfies the antecedent, but not the consequent of the conditional; and (we add this to Nicod's statement) it is neutral, or irrelevant, with respect to the hypothesis if it does not satisfy the antecedent.

This criterion can readily be extended so as to be applicable also to universal conditionals containing more than one quantifier, such as 'Twins always resemble each other', or, in symbolic notation, '$(x)(y)(\text{Twins}(x, y) \supset \text{Rsbl}(x, y))$'. In these cases, a confirming instance consists of an ordered couple, or triple, etc., of objects satisfying the antecedent and the consequent of the conditional. (In the case of the last illustration, any two persons who are twins and resemble each other would confirm the hypothesis; twins who do not resemble each other would disconfirm it; and any two persons not twins – no matter whether they resemble each other or not – would constitute irrelevant evidence.)

We shall refer to this criterion as Nicod's criterion.[4] It states explicitly what is perhaps the most common tacit interpretation of the concept of confirmation. While seemingly quite adequate, it suffers from serious shortcomings, as will now be shown.

(*a*) First, the applicability of this criterion is restricted to hypotheses of universal conditional form; it provides no standards of confirmation for existential hypotheses (such as 'There exists organic life on other stars', or 'Poliomyelitis is caused by some virus') or for hypotheses whose explicit formulation calls for the use of both universal and existential quantifiers (such as 'Every human being dies some finite number of years after his birth', or the psychological hypothesis, 'You can fool all of the people some of the time and some of the people all of the time, but you cannot fool all of the people all of the time', which may be symbolized by '$(x)(Et)\text{Fl}(x, t) \cdot (Ex)(t)\text{Fl}(x, t) \cdot \sim (x)(t)\text{Fl}(x, t)$', (where '$\text{Fl}(x, t)$' stands for 'You can fool person x at time t'). We note, therefore, the desideratum of establishing a criterion of confirmation which is applicable to hypotheses of *any* form.[5]

(*b*) We now turn to a second shortcoming of Nicod's criterion. Consider the two sentences

S_1: '$(x)[\text{Raven}(x) \supset \text{Black}(x)]$';
S_2: '$(x)[\sim\text{Black}(x) \supset \sim \text{Raven}(x)]$'

(i.e. 'All ravens are black' and 'Whatever is not black is not a raven'), and let a, b, c, d be four objects such that a is a raven and black, b a raven but not black, c not a raven but black, and d neither a raven nor black. Then according to Nicod's criterion, a would confirm S_1, but be neutral with respect to S_2; b would disconfirm both S_1 and S_2; c would be neutral with respect to both S_1 and S_2, and d would confirm S_2, but be neutral with respect to S_1.

But S_1 and S_2 are logically equivalent; they have the same content, they are different formulations of the same hypothesis. And yet, by Nicod's criterion, either of the objects a and d would be confirming for one of the two sentences, but neutral with respect to the other. This means that Nicod's criterion makes confirmation depend not only on the content of the hypothesis, but also on its formulation.[6]

One remarkable consequence of this situation is that every hypothesis to which the criterion is applicable – i.e. every universal conditional – can be stated in a form for which there cannot possibly exist any confirming instances. Thus, e.g. the sentence

$$(x)[(\text{Raven}(x) \cdot \sim \text{Black}(x)) \supset (\text{Raven}(x) \cdot \sim \text{Raven}(x)]$$

is readily recognized as equivalent to both S_1 and S_2 above; yet no object whatever can confirm this sentence, i.e. satisfy both its antecedent and its consequent; for the consequent is contradictory. An analogous transformation is, of course, applicable to any other sentence of universal conditional form.

4 The Equivalence Condition

The results just obtained call attention to the following condition which an adequately defined concept of confirmation should satisfy, and in the light of which Nicod's criterion has to be rejected as inadequate:

Equivalence condition: Whatever confirms (disconfirms) one of two equivalent sentences, also confirms (disconfirms) the other.

Fulfillment of this condition makes the confirmation of a hypothesis independent of the way in which it is formulated; and no doubt it will be conceded that this is a necessary condition for the adequacy of any proposed criterion of confirmation. Otherwise, the question as to whether certain data confirm a given hypothesis would have to be answered by saying: "That depends on which of the different equivalent formulations of the hypothesis is considered" – which appears absurd. Furthermore – and this is a more important point than an appeal to a feeling of absurdity – an adequate definition of confirmation will have to do justice to the way in which empirical hypotheses function in theoretical scientific contexts such as explanations and predictions; but when hypotheses are used for purposes of explanation or prediction,[7] they serve as premises in a deductive argument whose conclusion is a description of the event to be explained or predicted. The deduction is governed by the principles of formal logic, and according to the latter, a deduction which is valid will remain so if some or all of the premises are replaced by different but equivalent statements; and indeed, a scientist will feel free, in any theoretical reasoning involving certain hypotheses, to use the latter in whichever of their equivalent formulations are most convenient for the development of his conclusions. But if we adopted a concept of confirmation which did not satisfy the equivalence condition, then it would be possible, and indeed necessary, to argue in certain cases that it was sound scientific procedure to base a prediction on a given hypothesis if formulated in a sentence S_1, because a good deal of confirming evidence had been found for S_1; but that it was altogether inadmissible to base the prediction (say, for convenience of deduction) on an equivalent formulation S_2, because no confirming evidence for S_2 was available. Thus, the equivalence condition has to be regarded as a necessary condition for the adequacy of any definition of confirmation.

5 The Paradoxes of Confirmation

Perhaps we seem to have been laboring the obvious in stressing the necessity of satisfying the equivalence condition. This impression is likely to vanish upon consideration of certain consequences which derive from a combination of the equivalence condition with a most natural and plausible assumption concerning a sufficient condition of confirmation.

The essence of the criticism we have leveled so far against Nicod's criterion is that it certainly cannot serve as a necessary condition of confirmation; thus, in the illustration given in the beginning of section 3, object *a* confirms S_1 and should therefore also be considered as confirming S_2, while according to Nicod's criterion it is not. Satisfaction of the latter is therefore not a necessary condition for confirming evidence.

On the other hand, Nicod's criterion might still be considered as stating a particularly obvious and important sufficient condition of confirmation. And indeed, if we restrict ourselves to universal conditional hypotheses in one variable[8] – such as S_1 and S_2 in the above illustration – then it seems perfectly reasonable to qualify an object as confirming such a hypothesis if it satisfies both its antecedent and its consequent. The plausibility of this view will be further corroborated in the course of our subsequent analyses.

Thus, we shall agree that if *a* is both a raven and black, then *a* certainly confirms S_1: '$(x)(\text{Raven}(x) \supset \text{Black}(x))$', and if *d* is neither black nor a raven, *d* certainly confirms S_2: $(x)[\sim \text{Black}(x) \supset \sim \text{Raven}(x)]$'.

Let us now combine this simple stipulation with the equivalence condition. Since S_1 and S_2 are equivalent, *d* is confirming also for S_1; and thus, we have to recognize as confirming for S_1 any object which is neither black nor a raven. Consequently, any red pencil, any green leaf, any yellow cow, etc., becomes confirming evidence for the hypothesis that all ravens are black. This surprising consequence of two very adequate assumptions (the equivalence condition and the above sufficient condition of confirmation) can be further expanded: The sentence S_1 can readily be shown to be equivalent to S_3: '$(x)[(\text{Raven}(x) \vee \sim \text{Raven}(x)) \supset (\sim \text{Raven}(x) \vee \text{Black}(x))]$', i.e. 'Anything which is or is not a raven is either no raven or black'. According to the above sufficient condition, S_3 is certainly confirmed by any object, say *e*, such that (1) *e* is or is not a raven and, in addition (2) *e* is not a raven or is also black. Since (1) is analytic, these conditions reduce to (2). By virtue of the equivalence condition, we have therefore to consider as confirming for S_1 any object which is either no raven or also black (in other words: any object which is no raven at all, or a black raven).

Of the four objects characterized in section 3, *a*, *c* and *d* would therefore constitute confirming evidence for S_1, while *b* would be disconforming for S_1. This implies that any nonraven represents confirming evidence for the hypothesis that all ravens are black.[9]

We shall refer to these implications of the equivalence condition and of the above sufficient condition of confirmation as the *paradoxes of confirmation*.

How are these paradoxes to be dealt with? Renouncing the equivalence condition would not represent an acceptable solution, as it is shown by the considerations presented in section 4. Nor does it seem possible to dispense with the stipulation that an object satisfying two conditions, C_1 and C_2, should be considered as confirming a general hypothesis to the effect that any object which satisfies C_1 also satisfies C_2.

But the deduction of the above paradoxical results rests on one other assumption which is usually taken for granted, namely, that the meaning of general empirical hypotheses, such as that all ravens are black, or that all sodium salts burn yellow, can be adequately expressed by means of sentences of universal conditional form, such as '$(x)[\text{Raven}(x) \supset \text{Black}(x)]$' and '$(x)(\text{Sod. Salt}(x) \supset \text{Burn Yellow}(x))$', etc. Perhaps this customary mode of presentation has to be modified; and perhaps such a modification would automatically remove the paradoxes of confirmation? If this is not so, there seems to be only one alternative left, namely to show that the impression of the paradoxical character of those consequences is due to misunderstanding and can be dispelled, so that no theoretical difficulty remains. We shall now consider these two possibilities in turn: Subsections 5.11 and 5.12 are devoted to a discussion of two different proposals for a modified representation of general hypotheses; in subsection 5.2, we shall discuss the second alternative, i.e. the possibility of tracing the impression of paradoxicality back to a misunderstanding.

5.11. It has often been pointed out that while Aristotelian logic, in agreement with prevalent everyday usage, confers existential import upon sentences of the form 'All *P*'s are *Q*'s', a universal conditional sentence, in the sense of modern logic, has no existential import; thus, the sentence

$$\text{'}(x)[\text{Mermaid}(x) \supset \text{Green}(x)]\text{'}$$

does not imply the existence of mermaids; it merely asserts that any object either is not a mermaid at all, or a green mermaid; and it is true simply because of the fact that there are no mermaids. General laws and hypotheses in science, however – so it might be argued – are meant to have existential import; and one might attempt to express the latter by supplementing the customary universal conditional by an existential clause. Thus, the hypothesis that all ravens are black would be expressed by means of the sentence S_1: '$[(x)(\text{Raven}(x) \supset \text{Black}(x)] \cdot (Ex)\text{Raven}(x)$'; and the hypothesis that no nonblack things are ravens by S_2: '$(x)[\sim\text{Black}(x) \supset \sim \text{Raven}(x)] \cdot (Ex) \sim \text{Black}(x)$'. Clearly, these sentences are not equivalent, and of the four objects *a*, *b*, *c*, *d* characterized in section 3, part (*b*), only *a* might reasonably be said to confirm S_1, and only *d* to confirm S_2. Yet this method of avoiding the paradoxes of confirmation is open to serious objections:

(*a*) First of all, the representation of every general hypothesis by a conjunction of a universal conditional and an existential sentence would invalidate many logical inferences which are generally accepted as permissible in a theoretical argument. Thus, for example, the assertions that all sodium salts burn yellow, and that whatever does not burn yellow is no sodium salt are logically equivalent according to customary understanding and usage, and their representation by universal conditionals preserves this equivalence; but if existential clauses are added, the two assertions are no longer equivalent, as is illustrated above by the analogous case of S_1 and S_2.

(*b*) Second, the customary formulation of general hypotheses in empirical science clearly does not contain an existential clause, nor does it, as a rule, even indirectly determine such a clause unambiguously. Thus, consider the hypothesis that if a person after receiving an injection of a certain test substance has a positive skin reaction, he has diphtheria. Should we construe the existential clause here as referring to persons, to persons receiving the injection, or to persons who, upon receiving the injection, show a positive skin reaction? A more or less arbitrary decision has to be made; each of the possible decisions gives a different interpretation to the hypothesis, and none of them seems to be really implied by the latter.

(c) Finally, many universal hypotheses cannot be said to imply an existential clause at all. Thus, it may happen that from a certain astrophysical theory a universal hypothesis is deduced concerning the character of the phenomena which would take place under certain specified extreme conditions. A hypothesis of this kind need not (and, as a rule, does not) imply that such extreme conditions ever were or will be realized; it has no existential import. Or consider a biological hypothesis to the effect that whenever man and ape are crossed, the offspring will have such and such characteristics. This is a general hypothesis; it might be contemplated as a mere conjecture, or as a consequence of a broader genetic theory, other implications of which may already have been tested with positive results; but unquestionably the hypothesis does not imply an existential clause asserting that the contemplated kind of cross-breeding referred to will, at some time, actually take place.

5.12 Perhaps the impression of the paradoxical character of the cases discussed in the beginning of section 5 may be said to grow out of the feeling that the hypothesis that all ravens are black is about ravens, and not about nonblack things, nor about all things. The use of an existential clause was one attempt at exhibiting this presumed peculiarity of the hypothesis. The attempt has failed, and if we wish to express the point in question, we shall have to look for a stronger device. The idea suggests itself of representing a general hypothesis by the customary universal conditional, supplemented by the indication of the specific "field of application" of the hypothesis; thus, we might represent the hypothesis that all ravens are black by the sentence '(x) [Raven$(x) \supset$ Black(x)]' or any one of its equivalents, plus the indication 'Class of ravens', characterizing the field of application; and we might then require that every confirming instance should belong to the field of application. This procedure would exclude the objects *c* and *d* from those constituting confirming evidence and would thus avoid those undesirable consequences of the existential-clause device which were pointed out in 5.11 (*c*). But apart from this advantage, the second method is open to objections similar to those which apply to the first: (*a*) The way in which general hypotheses are used in science never involves the statement of a field of application; and the choice of the latter in a symbolic formulation of a given hypothesis thus introduces again a considerable measure of arbitrariness. In particular, for a scientific hypothesis to the effect that all *P*'s are *Q*'s, the field of application cannot simply be said to be the class of all *P*'s; for a hypothesis such as that all sodium salts burn yellow finds important application in tests with negative results; e.g., it may be applied to a substance of which it is not known whether it contains sodium salts, nor whether it burns yellow; and if the flame does not turn yellow, the hypothesis serves to establish the absence of sodium salts. The same is true of all other hypotheses used for tests of this type. (*b*) Again, the consistent use of a field of application in the formulation of general hypotheses would involve considerable logical complications, and yet would have no counterpart in the theoretical procedure of science, where hypotheses are subjected to various kinds of logical transformation and inference without any consideration that might be regarded as referring to changes in the fields of application. This method of meeting the paradoxes would therefore amount to dodging the problem by means of an ad hoc device which cannot be justified by reference to actual scientific procedure.

5.2 We have examined two alternatives to the customary method of representing general hypotheses by means of universal conditionals; neither of them proved an adequate means of precluding the paradoxes of confirmation. We shall now try to show that what is wrong does not lie in the customary way of construing and representing general hypotheses, but rather in our reliance on a misleading intuition in the matter: The impression of a paradoxical situation is not objectively founded; it is a psychological illusion.

(*a*) One source of misunderstanding is the view, referred to before, that a hypothesis of the simple form 'Every P is a Q', such as 'All sodium salts burn yellow', asserts something about a certain limited class of objects only, namely, the class of all P's. This idea involves a confusion of logical and practical considerations: Our interest in the hypothesis may be focussed upon its applicability to that particular class of objects, but the hypothesis nevertheless asserts something about, and indeed imposes restrictions upon, *all* objects (within the logical type of the variable occurring in the hypothesis, which in the case of our last illustration might be the class of all physical objects). Indeed, a hypothesis of the form 'Every P is a Q' forbids the occurrence of any objects having the property P but lacking the property Q; i.e. it restricts all objects whatsoever to the class of those which either lack the property P or also have the property Q. Now, every object either belongs to this class or falls outside it, and thus, every object – and not only the P's – either conforms to the hypothesis or violates it; there is no object which is not implicitly referred to by a hypothesis of this type. In particular, every object which either is no sodium salt or burns yellow conforms to, and thus bears out, the hypothesis that all sodium salts burn yellow; every other object violates that hypothesis.

The weakness of the idea under consideration is evidenced also by the observation that the class of objects about which a hypothesis is supposed to assert something is in no way clearly determined, and that it changes with the context, as was shown in 5.12 (*a*).

(*b*) A second important source of the appearance of paradoxicality in certain cases of confirmation is exhibited by the following consideration.

Suppose that in support of the assertion 'All sodium salts burn yellow' somebody were to adduce an experiment in which a piece of pure ice was held into a colorless flame and did not turn the flame yellow. This result would confirm the assertion, 'Whatever does not burn yellow is no sodium salt' and consequently, by virtue of the equivalence condition, it would confirm the original formulation. Why does this impress us as paradoxical? The reason becomes clear when we compare the previous situation with the case where an object whose chemical constitution is as yet unknown to us is held into a flame and fails to turn it yellow, and where subsequent analysis reveals it to contain no sodium salt. This outcome, we should no doubt agree, is what was to be expected on the basis of the hypothesis that all sodium salts burn yellow – no matter in which of its various equivalent formulations it may be expressed; thus, the data here obtained constitute confirming evidence for the hypothesis. Now the only difference between the two situations here considered is that in the first case we are told beforehand the test substance is ice, and we happen to "know anyhow" that ice contains no sodium salt; this has the consequence that the outcome of the flame-color test becomes entirely irrelevant for the confirmation of the hypothesis and thus can yield no new evidence for us. Indeed, if the flame should not turn yellow, the hypothesis requires that the substance contain no sodium salt – and we know beforehand that ice does not; and if the flame should turn yellow, the hypothesis would impose no further restrictions on the substance: hence, either of the possible outcomes of the experiment would be in accord with the hypothesis.

The analysis of this example illustrates a general point: In the seemingly paradoxical cases of confirmation, we are often not actually judging the relation of the given evidence E alone to the hypothesis H (we fail to observe the methodological fiction, characteristic of every case of confirmation, that we have no relevant evidence for H other than that included in E); instead, we tacitly introduce a comparison of H with a body of evidence which consists of E in conjunction with additional information that we happen to have at our disposal; in our illustration, this information includes the knowledge (1) that the substance used in the experiment is ice, and (2) that ice contains no sodium salt. If we assume this additional information as given, then, of course, the outcome of the experiment can add no strength to the hypothesis under consideration. But if we are careful to avoid this tacit reference to additional knowledge (which entirely changes the character of the problem), and if we formulate the question as to the confirming character of the evidence in a manner adequate to the concept of confirmation as used in this paper, we have to ask: Given some object *a* (it happens to be a piece of ice, but this fact is not included in the evidence), and given the fact that *a* does not turn the flame yellow and is no sodium salt: does *a* then constitute confirming evidence for the hypothesis? And now – no matter whether *a* is ice or some other substance – it is clear that the answer has to be in the affirmative; and the paradoxes vanish.

So far, in section (*b*), we have considered mainly that type of paradoxical case which is illustrated by the assertion that any nonblack nonraven constitutes confirming evidence for the hypothesis, 'All ravens are black.' However, the general idea just outlined applies as well to the even more extreme cases exemplified by the assertion that any nonraven as well as any black object confirms the hypothesis in question. Let us illustrate this by reference to the latter case. If the given evidence *E* – i.e. in the sense of the required methodological fiction, all data relevant for the hypothesis – consists only of one object which, in addition, is black, then *E* may reasonably be said to support even the hypothesis that all objects are black, and a fortiori *E* supports the weaker assertion that all ravens are black. In this case, again, our factual knowledge that not all objects are black tends to create an impression of paradoxicality which is not justified on logical grounds. Other paradoxical cases of confirmation may be dealt with analogously. Thus it turns out that the paradoxes of confirmation, as formulated above, are due to a misguided intuition in the matter rather than to a logical flaw in the two stipulations from which they were derived.[10,11]

Notes

1 Jean Nicod, *Foundations of Geometry and Induction* (trans. by P. P. Wiener), London, 1930; 219; *cf.* also R. M. Eaton's discussion of "Confirmation and Infirmation," which is based on Nicod's views; it is included in Chap. III of his *General Logic* (New York, 1931).

2 In this essay, only the most elementary devices of this notation are used; the symbolism is essentially that of *Principia Mathematica*, except that parentheses are used instead of dots, and that existential quantification is symbolized by '(E)' instead of by the inverted 'E.'

3 (Added in 1964). More precisely we would have to say, in Nicod's parlance, that the hypothesis is confirmed by the *proposition* that *a* is both *P* and *Q*, and is disconfirmed by the *proposition* that *a* is *P* but not *Q*.

4 This term is chosen for convenience, and in view of the above explicit formulation given by Nicod; it is not, of course, intended to imply that this conception of confirmation originated with Nicod.

5 For a rigorous formulation of the problem, it is necessary first to lay down assumptions as to the means of expression and the logical structure of the language in which the hypotheses are supposed to be formulated; the desideratum then calls for a definition of confirmation applicable to any hypothesis which can be expressed in the given language. Generally speaking, the problem becomes increasingly difficult with increasing richness and complexity of the assumed language of science.

6 This difficulty was pointed out, in substance, in my article "Le problème de la vérité," *Theoria* (Göteborg), vol. 3 (1937), esp. p. 222.

7 For a more detailed account of the logical structure of scientific explanation and prediction, *cf.* C. G. Hempel, "The Function of General Laws in History," *The Journal of Philosophy*, vol. 39 (1942), esp. sections 2, 3, 4. The characterization, given in that paper as well as in the above text, of explanations and predictions as arguments of a deductive logical structure, embodies an oversimplification: as will be shown in section 7 of the present essay, explanations and predictions often involve "quasi-inductive" steps besides deductive ones. This point, however, does not affect the validity of the above argument.

8 This restriction is essential: In its general form which applies to universal conditionals in any number of variables, Nicod's criterion cannot even be construed as expressing a sufficient condition of confirmation. This is shown by the following rather surprising example: Consider the hypothesis:

$$S_1: (x)(y)[\sim (R(x, y) \cdot R(y, x)) \supset (R(x, y) \cdot \sim R(y, x))]$$

Let *a*, *b* be two objects such that $R(a, b)$ and $\sim R(b, a)$. Then clearly, the couple (a, b) satisfies both the antecedent and the consequent of the universal conditional S_1; hence, if Nicod's criterion in its general form is accepted as stating a sufficient condition of confirmation, (a, b) constitutes confirming evidence for S_1. But S_1 can be shown to be equivalent to

$$S_2: (x)(y)R(x, y)$$

Now, by hypothesis, we have $\sim R(b, a)$; and this flatly contradicts S_2 and thus S_1. Thus, the couple (a, b), although satisfying both the antecedent and the consequent of the universal conditional S_1, actually constitutes disconfirming evidence of the strongest kind (conclusively disconfirming evidence, as we shall say later) for that sentence. This illustration reveals a striking and – as far as I am aware – hitherto unnoticed weakness of that conception of confirmation which underlies Nicod's criterion. In order to realize the bearing of our illustration upon Nicod's original formulation, let A and B be $\sim (R(x, y) \cdot R(y, x))$ and $R(x, y) \cdot \sim (R(y, x))$, respectively. Then S_1 asserts that A entails B, and the couple (a, b) is a case of the presence of B in the presence of A; this should, according to Nicod, be favorable to S_1.

9 (Added in 1964). The following further "paradoxical" consequence of our two conditions might be noted: Any hypothesis of universal conditional form can be equivalently rewritten as another hypothesis of the same form which, even if true, can have no confirming instances in Nicod's sense at all, since the proposition that a given object satisfies the antecedent and the consequent of the second hypothesis is self-contradictory. For example, '$(x)[P(x) \supset Q(x)]$' is equivalent to the sentence '$(x)[(P(x) \cdot \sim Q(x)) \supset (P(x) \cdot \sim P(x))]$', whose consequent is true of nothing.

10 The basic idea of section (*b*) in the above analysis is due to Dr Nelson Goodman, to whom I wish to reiterate my thanks for the help he rendered me, through many discussions, in clarifying my ideas on this point.

11 The considerations presented in section (*b*) above are also influenced by, though not identical in content with, the very illuminating discussion of the paradoxes by the Polish methodologist and logician Janina Hosiasson-Lindenbaum; *cf.* her article "On Confirmation," *The Journal of Symbolic Logic*, vol. 5 (1940), especially section 4. Dr Hosiasson's attention had been called to the paradoxes by my article "Le problème de la vérité" (*cf.* note 20) and by discussions with me. To my knowledge, hers has so far been the only publication which presents an explicit attempt to solve the problem. Her solution is based on a theory of degrees of confirmation, which is developed in the form of an uninterpreted axiomtaic system, and most of her arguments presuppose that theoretical framework. I have profited, however, by some of Miss Hosiasson's more general observations which proved relevant for the analysis of the paradoxes of the nongraduated or qualitative concept of confirmation which forms the object of the present study.

One point in those of Miss Hosiasson's comments which rest on her theory of degrees of confirmation is of particular interest, and I should like to discuss it briefly. Stated in reference to the raven hypothesis, it consists in the suggestion that the finding of one nonblack object which is no raven, while constituting confirming evidence for the hypothesis, would increase the degree of confirmation of the hypothesis by a smaller amount than the finding of one raven which is black. This is said to be so because the class of all ravens is much less numerous than that of all nonblack objects, so that – to put the idea in suggestive though somewhat misleading terms – the finding of one black raven confirms a larger portion of the total content of the hypothesis than the finding of one nonblack nonraven. In fact, from the basic assumptions of her theory, Miss Hosiasson is able to derive a theorem according to which the above statement about the relative increase in degree of confirmation will hold provided that actually the number of all ravens is small compared with the number of all nonblack objects. But is this last numerical assumption actually warranted in the present case and analogously in all other "paradoxical" cases? The answer depends in part upon the logical structure of the language of science. If a "coordinate language" is used, in which, say, finite space–time regions figure as individuals, then the raven hypothesis assumes some such form as 'Every space–time region which contains a raven contains something black'; and even if the total number of ravens ever to exist is finite, the class of space–time regions containing a raven has the power of the continuum, and so does the class of space–time regions containing something nonblack; thus, for a coordinate language of the type under consideration, the above numerical assumption is not warranted. Now the use of a coordinate language may appear quite artificial in this particular illustration; but it will seem very appropriate in many other contexts, such as, e.g., that of physical field theories. On the other hand, Miss Hosiasson's numerical assumption may well be justified on the basis of a "thing language," in which physical objects of finite size function as individuals. Of course, even on this basis, it remains an empirical question, for every hypothesis of the form 'All P's are Q's', whether actually the class of non-Q's is much more numerous than the class of P's; and in many cases this question will be very difficult to decide.