

HEMPEL AND THE PARADOXES OF CONFIRMATION

Jan Sprenger

1 TOWARDS A LOGIC OF CONFIRMATION

The beginning of modern philosophy of science is generally associated with the label of *logical empiricism*, in particular with the members of the Vienna Circle. Some of them, as Frank, Hahn and Neurath, were themselves scientists, others, as Carnap and Schlick, were philosophers, but deeply impressed by the scientific revolutions at the beginning of the 20th century. All of them were unified in admiration for the systematicity and enduring success of science. This affected their philosophical views and led to a sharp break with the “metaphysical” philosophical tradition and to a re-invention of empiricist epistemology with a strong emphasis on science, our best source of high-level knowledge. Indeed, the members of the Vienna Circle were scientifically trained and used to the scientific method of test and observation. For them, metaphysical claims were neither verifiable nor falsifiable through empirical methods, and therefore neither true nor false, but meaningless. Proper philosophical analysis had to separate senseless (metaphysical) from meaningful (empirical) claims and to investigate our most reliable source of knowledge: science.¹

The latter task included the development of formal frameworks for discovering the logic of scientific method and progress. Rudolf Carnap, who devoted much of his work to this task, was in fact one of the most influential figures of the Vienna Circle. In 1930, Carnap and the Berlin philosopher Hans Reichenbach took over the journal ‘Annalen der Philosophie’ and renamed it ‘Erkenntnis’. Under that name, it became a major publication organ for the works of the logical empiricists. The German-Austrian collaboration in the editorial board was no matter of chance: Congenial to the Vienna group, several similar-minded researchers based in Berlin gathered in the ‘Berlin Society for Empirical Philosophy’, among them Reichenbach. It was here that a young German student of mathematics, physics and philosophy — Carl Gustav Hempel — got into contact with empiricist philosophy. On the 1929 conference on the epistemology of the exact sciences in Berlin, he got to know Carnap and soon moved to Vienna himself. Nevertheless, he obtained his doctorate degree in Berlin in 1934, but faced with Nazi rule, Hempel

¹Cf. [Friedman, 1999; Uebel, 2006].

soon opted for emigration and later became Carnap's assistant at the University of Chicago. Thus it is no matter of chance that the contents of Carnap's and Hempel's philosophy are so close to each other. Similar to Carnap, Hempel was interested in the logic of science and in particular in the problem of inductive inference. Similar to Carnap, Hempel thought that the introduction of rigorous methods would help us to establish a logic of induction and confirmation. Carnap's life project consisted in developing a probabilistic logic of induction (Carnap [1950], cf. [Zabell, 2009]), similar to a deductive calculus for truth-preserving inferences. Indeed, the success of the calculus of deductive logic suggests a similar calculus for inductive, ampliative inferences that could be applied to the confirmation of scientific hypotheses. Having a logic of confirmation would thus contribute to the central aims of empiricist philosophy: to understand the progress and success of science, and in particular the replacement of old by new theories and the testability of abstract hypotheses by empirical observations. While in principle cherishing Carnap's probabilistic work in that area, Hempel had some subtle methodological reservations: Prior to explicating the concept of confirmation in a probabilistic framework, we are supposed to clarify our *qualitative* concept of confirmation and to develop general adequacy criteria for an explication of confirmation. Therefore my essay deals less with probabilistic than with qualitative approaches to confirmation theory in modern philosophy of science. Hempel's main contribution, the essay 'Studies in the Logic of Confirmation', was published in 1945, right after first pioneer works in the area (e.g. [Hosiasson-Lindenbaum, 1940]), but before Carnap's [1950; 1952] major monographs. On the way, we will also stumble over Hempel's famous *paradoxes of confirmation*, which pose, or so I will argue, a great challenge for any account of confirmation.²

Let us begin with some preliminary thoughts. In science, confirmation becomes an issue whenever science interacts with the world, especially when scientific hypotheses are subjected to empirical tests. Where exactly can a logic of induction and confirmation help us? Hempel distinguishes three stages of empirical testing [Hempel, 1945/1965, pp. 40-41]: First, we design, set up and carefully conduct scientific experiments, we try to avoid misleading observations, double-check the data, clear them up and finally bring them into a canonical form that we can use in

²Confirmation is generally thought to hold between a hypothesis and pieces of evidence — a piece of evidence proves, confirms, undermines, refutes or is irrelevant to a hypothesis. At first sight, it sounds plausible to think of confirmation as a *semantic* relation between a scientific theory on the one side and a real-world object on the other side. For instance, a black raven seems to confirm the hypothesis that all ravens are black. But recall that we would like to assimilate confirmation theory to deductive logic and to find a system of *syntactic* rules for valid inductive inference. Therefore we should frame the evidence into sentences of a (formal) language, in order to gain access to powerful logical tools, e.g. checking deducibility and consistency relations between evidence and hypothesis. Thus, Hempel argues, a purely semantic account of confirmation is inadequate. We should set up a syntactic relation between hypothesis and evidence where both relata are (sets of) first-order sentences. (Cf. [Hempel, 1945/1965, pp. 21-22]). When I nevertheless say that 'a black raven confirms hypothesis H ', this is just a matter of convenience and means the corresponding observation report 'there is a black raven'.

the next stage.³ In the second stage, these data are brought to bear on the hypothesis at stake — do they constitute supporting or undermining evidence? Third and last, the hypothesis is re-assessed on the basis of a judgment of confirmation or disconfirmation: we decide to accept it, to reject it or to suspend judgment and to collect further evidence. — In these three stages, only the second stage is, or so Hempel argues, accessible to a logical analysis: the first and the third stage are full of pragmatically loaded decisions, e.g. which experiment to conduct, how to screen off the data against external nuisance factors, or which strength of evidence is required for accepting a hypothesis. Evidently, those processes cannot be represented by purely formal means. That's different for the second stage which compares observational sentences (in which the evidence is framed) with theoretical sentences which represent the hypothesis or theory. This is the point where logical tools can help to analyze the relation between both kinds of sentences and to set up criteria for successful scientific confirmation.

A fundamental objection against a logic of confirmation holds that scientists frequently disagree whether an empirical finding really confirms a theoretical hypothesis, and this phenomenon is too common to ascribe it to irrationality on behalf of the researchers. Common scientific sense may not be able to decide such questions, first because the case under scrutiny might be very complicated and second, because people might have different ideas of common sense in a specific case. Formal criteria of confirmation help to settle the discussion, and once again, it is helpful to consider the analogy to deductive logic. For each valid deductive inference, there is a deduction of the conclusion from the logical axioms (that is the completeness theorem for first-order logic). Hence, in case there is a disagreement about the validity of a deductive inference, the formal tools can help us to settle the question. In the same way that the validity of a deductive inference can be checked using formal tools (deductions), it is desirable to have formal tools which examine the validity of an inductive inference. Sometimes this project is deemed futile because scientists do not always make their criteria of confirmation explicit. But that objection conflates a logical with a psychological point [Hempel, 1945/1965, pp. 9-10] — the lack of explicit confirmation criteria in scientific practice does not refute their existence. The objection merely shows that if such criteria exist, scientists are often not aware of them. But since scientists make, in spite of all disagreement in special cases, in general consistent judgments on evidential relevance, this is still a fruitful project. Confirmation theory thus aims at a rational reconstruction of inductive practice that is not only descriptively adequate, but also able to correct methodological mistakes in science. Thus confirmation theory is vastly more than a remote philosophical subdiscipline, it is actually a proper part of the foundations of science, in the very spirit of logical empiricism. Later *œuvres* where debates about proper scientific method interfere with confirmation-theoretic problems (e.g. [Royall, 1997]) vindicate this view. Let us now review Hempel's pioneer work.

³Suppes [1962/1969] refers to this activity as building 'models of data'.

2 ADEQUACY CRITERIA

Carnap and Hempel both worked on an explication of confirmation, but their methods were quite different. While Carnap connected confirmation to probability by proposing ‘degree of confirmation’ as an interpretation of probability, Hempel pursued a non-probabilistic approach which precedes the quantitative analysis. His method can be described thus: At the beginning, general considerations yield adequacy criteria for every sensible account of confirmation [Hempel, 1945/1965, pp. 30-33], considerably narrowing down the space of admissible accounts. Out of the remaining accounts, Hempel selects the one that also captures a core intuition about confirmation, namely that hypotheses are confirmed by their instances. Let us now see which criteria Hempel develops. The first criterion which he suggests is the

Entailment Condition (EnC): If the observation report E logically implies the hypothesis H then E confirms H .

For example, if the hypothesis reads ‘there are white ravens’ then, obviously, the observation of a white raven *proves* it and *a fortiori*, confirms it: Logical implication is the strongest possible form of evidential support. So the Entailment Condition sounds very reasonable.

Then, if a theory is confirmed by a piece of evidence, it seems strange to deny that consequences of the theory are not confirmed by the evidence. For instance, if observations confirm Newton’s law of gravitation, they should confirm Kepler’s laws, too, since the latter’s predictions have to agree with the gravitation law. In other words, we demand satisfaction of the

Consequence Condition (CC): If an observation report E confirms every member of a set of sentences \mathcal{S} , then it confirms every consequence of \mathcal{S} (e.g. every sentence H for which $\mathcal{S} \models H$).

In fact, the consequence condition is quite powerful, and several natural adequacy criteria follow from it. For instance, the

Equivalence Condition (EC): If H and H' are logically equivalent sentences, then the observation report E confirms H if and only if E confirms H' .⁴

It is straightforward to see that (EC) follows from (CC): If a sentence H is confirmed by E and H is equivalent to H' , then H' is a logical consequence of $\{H\}$ and the Consequence Condition can be applied, yielding that H' is also confirmed by E , and vice versa.

⁴This condition can naturally be extended to a condition for the evidence, asserting that the confirmation relation is invariant under replacing the evidence statement by logically equivalent statements.

Certainly, the equivalence condition is a minimal constraint on any account of confirmation. We have already said that scientific hypotheses are usually framed in the logical vocabulary of first-order logic (or a reduct thereof). That allows us to state them in different, but logically equivalent forms.⁵ The idea of the equivalence condition is that ‘saying the same with different words’ does not make a difference with regard to relations of confirmation and support: Hypotheses which express the same content in different words are equally supported and undermined by a piece of evidence, independent of the chosen formulation. To see this in more detail, note that for deductive relations, the Equivalence Condition holds by definition: If A logically implies B , A also implies any B' that is logically equivalent to B . An account of confirmation should contain relations of deduction and entailment as special cases: If an observation entailed the negation of a hypothesis, in other words, if the hypothesis were *falsified* by actual evidence, this would equally speak against all equivalent versions and formulations of that hypothesis. Deduction and logical entailment do not make a difference between equivalent sentences, and logical and mathematical axiomatizations are typical of the modern exact sciences (e.g. the propagation of sound is described by a general theory of mechanic waves). If the Equivalence Condition did not hold, the degree of support which a hypothesis got would depend on the specific formulation of the hypothesis. But that would run counter to all efforts to introduce exact mathematical methods into science, thereby making scientific analysis more precise, and ultimately more successful.

Obviously, the Consequence Condition also implies the

Special Consequence Condition (SCC): If an observation report E confirms a hypothesis H , then it confirms every consequence of H .

However, there is an important confirmation intuition that contradicts (SCC) and stems from the link between prediction, test and confirmation. When a theory makes a prediction and this prediction is indeed observed, those observations lend empirical support to the theory. Abstract theories, like the General Theory of Relativity (GTR), are often not directly testable. We have to focus on parts of them and to use those parts for deriving observational consequences. This agrees with falsificationist methodology (e.g. [Popper, 1963]) — we derive conjectures and predictions from a theory and test them versus the empirical world. For instance, Eddington’s observations of the solar eclipse in 1919 did not *prove* GTR, but merely confirmed one of its predictions — namely the bending of light by massive bodies. Had the outcome been different, GTR (or one of the auxiliary assumptions) would have been falsified.

Evidently, the stronger a theory, the higher its predictive power. In particular, if the theory T predicts the observation sentence E , E is also a prediction of any stronger theory T' . This line of reasoning suggests the

Converse Consequence Condition (CCC): If an observation report E confirms a hypothesis H , then it confirms every hypothesis H'

⁵For instance, the definition of compactness for sets of real numbers can be stated in topological or in analytical terms.

that logically implies H (i.e. $H' \models H$).

Obviously, the Converse Consequence Condition (CCC) stands in sharp contrast to the Special Consequence Condition (SCC). Indeed, accepting both adequacy conditions at once would trivialize the concept of confirmation: Every observation report E trivially implies itself, so by (EnC), E confirms E . By (CCC), E also confirms $E.H$ for any hypothesis H since $E.H$ logically implies H . Since $E.H$ implies H and is confirmed by E , E confirms H by (SCC). Note that this derivation holds for an arbitrary hypothesis H and arbitrary observations E ! Our paradoxical result reveals that we have to make a decision between the prediction/observation-based scheme of inference (CCC) and the ‘conservative’ (SCC).

Hempel believed that the idea of predictive confirmation expressed in (CCC) is not an adequate image of confirmation in science. Sure, general laws as the law of gravitation are tested by observable consequences, such as the planetary motions. Indeed, successful tests of Kepler’s three laws are also believed to support the law of gravitation. But the evidence transfers from Kepler’s laws to the gravitation law because it is also an *instance* of the gravitation law — and not because the law of gravitation is logically stronger than Kepler’s laws. For instance, even the hypothesis ‘There is extraterrestrial life and Kepler’s laws hold’ is logically stronger than Kepler’s laws alone, but we would not like to say that this hypothesis can be confirmed by, let’s say, observing the orbit of Jupiter. If (CCC) is accepted, any hypothesis whatsoever (X) can be tacked to the confirmed hypothesis (H), and the new hypothesis $H.X$ is still confirmed by the evidence. These are the paradoxes of hypothetico-deductive confirmation, the *tacking paradoxes*.⁶ Moreover, (CCC) licenses the confirmation of mutually incompatible hypotheses: Let H be confirmed by E . Then both $H.X$ and $H.\neg X$ are, according to (CCC), confirmed by E . This sounds strange and arbitrary — the content of X is not at all relevant to our case, and if both hypotheses ($H.X$ and $H.\neg X$) are equally confirmed, it is not clear what we should *believe* in the end.

There are now two ways to proceed: Either we can try to restrict (CCC) to logically stronger hypotheses that stand in a relevance relation to the evidence. Then, the paradoxes vanish. Several authors have tried to mitigate the paradoxes of hypothetico-deductive confirmation along these lines, namely by the additional requirements that the tacked hypothesis $H.X$ or $H.\neg X$ be a *content part* of the hypothesis [Gemes, 1993] or that the inference to the evidence be ‘premise-relevant’ [Schurz, 1991]. So arbitrary hypotheses are no longer confirmed together with H . Hempel, however, chooses the other way — he rejects (CCC) in favor of (SCC). Contradictory hypotheses should, or so he argues, not be confirmed by one and the same evidence, in opposition to (CCC). We can put this view into another adequacy condition: hypotheses confirmed by a piece of evidence E must be consistent with each other.

Consistency Condition (CnC): If an observation report E confirms

⁶Cf. [Musgrave, 2009; Weisberg, 2009].

the hypotheses H and H' , then H is logically consistent with H' (i.e. there is at least one model of H that is also a model of H').

Finally, we summarize the three conditions that are essential to Hempel's account:

1. Entailment Condition (EnC): If $E \models H$, then E confirms H .
2. Consequence Condition (CC): If E confirms \mathcal{S} and $\mathcal{S} \models H$, then E confirms H . (Note: (CC) contains the Equivalence Condition (EC) and the Special Consequence Condition (SCC) as special cases.)
3. Consistency Condition (CnC): If E confirms H and H' , then H is logically consistent with H' .

3 THE SATISFACTION CRITERION

What should we demand of a piece of evidence in order to confirm a hypothesis? In general, logical entailment between evidence and hypothesis is too strong as a necessary criterion for confirmation. In particular, if the hypothesis is a universal conditional, no finite set of observations will ever be able to prove the hypothesis. But the evidence should certainly agree with those parts of the hypothesis that it is able to verify. Hempel suggests that, if an observation report says something about the singular terms a , b and c , the claims a hypothesis makes about a , b and c should be *satisfied* by the evidence. From such an observation report we could conclude that the hypothesis is true of the class of objects that occur in E . That is all we can demand of an confirming observation report, or so Hempel argues. In other words, we gain *instances* of a hypothesis from the evidence, and such instances confirm the hypothesis. To make this informal idea more precise, we have to introduce some definitions (partly taken from [Gemes, 2006]):

DEFINITION 1. An atomic well-formed formula (wff) β is *relevant to a wff* α if and only if there is some model M of α such that: if M' differs from M only in the value β is assigned, M' is not a model of α .

So intuitively, β is relevant for α if at least in one model of α the truth value of β cannot be changed without making α false. Now we can define the *domain* (or *scope*) of a wff:

DEFINITION 2. The *domain of a well-formed formula* α , denoted by $dom(\alpha)$, is the set of singular terms which occur in the atomic (!) well-formed formulas (wffs) of L that are relevant for α .

For example, the domain of $Fa.Fb$ is $\{a, b\}$ whereas the domain of $Fa.Ga$ is $\{a\}$ and the domain of $\forall x : Fx$ are all singular terms of the logical language. In other words, quantifiers are treated substitutionally. The domain of a formula is

thus the set of singular terms about which something is asserted. Those singular terms are said to occur essentially in the formula:

DEFINITION 3. A singular term a occurs essentially in a formula β if and only if a is in the domain of β .

So, i.e. a occurs essentially in $Fa.Fb$, but not in $(Fa \vee \neg Fa).Fb$. Now, we are interested in the development of a formula for the domain of a certain formula.

DEFINITION 4. The development of a formula H for a formula E , $H|_E$, is the restriction of H to the domain of E , i.e. to all singular terms that occur essentially in E .⁷

For instance, $(\forall x : Fx)|_{\{a,b\}}$ is $Fa.Fb$, and the development of the formula $\forall x : Fx$ for $Fa.Ga.Gb$ is $Fa.Fb$. Now we have the technical prerequisites for understanding Hempel's satisfaction criterion: The evidence entails the hypothesis not directly, but it entails the restriction of the hypothesis to the domain of the evidence.

DEFINITION 5. (Satisfaction criterion) A piece of evidence E directly Hempel-confirms a hypothesis H if and only if E entails the development of H to the domain of E . In other words, $E \models H|_{dom(E)}$.

DEFINITION 6. (Hempel-confirmation) A piece of evidence E Hempel-confirms a hypothesis H if and only if H is entailed by a set of sentences Γ so that for all sentences $\phi \in \Gamma$, ϕ is directly Hempel-confirmed by E .

There are also formulations of those criteria that refer to a body of knowledge which provides the background for evaluating the confirmation relation (e.g. our current theory of physics). We do not need that for illustrating Hempel's basic idea, but background information plays a crucial role in contrasting a hypothesis with empirical observations, as illustrated by the Duhem problem⁸: Does the failure of a scientific test speak against the hypothesis or against the auxiliary assumptions which we need for connecting the evidence to the hypothesis? Therefore we give a formulation of Hempel's satisfaction criterion which includes background knowledge.

DEFINITION 7. (Satisfaction criterion, triadic formulation) A piece of evidence E directly Hempel-confirms a hypothesis H relative to background knowledge K if and only if E and K jointly entail the development of H to the domain of E . In other words, $E.K \models H|_{dom(E)}$.⁹

DEFINITION 8. (Hempel-confirmation, triadic formulation) A piece of evidence E Hempel-confirms a hypothesis H relative to K if and only if H is entailed by a set of sentences Γ so that for all sentences $\phi \in \Gamma$, ϕ is directly Hempel-confirmed by E relative to K .

⁷The development of a formula can be defined precisely by a recursive definition, cf. [Hempel, 1943]. For our purposes, the informal version is sufficient.

⁸Cf. [Duhem, 1914].

⁹Cf. [Hempel, 1945/1965, pp. 36-37].

For example, Fa (directly) Hempel-confirms the hypothesis $\forall x : Fx$. Obviously, every piece of evidence that directly-Hempel confirms a hypothesis also Hempel-confirms it, but not vice versa.

It is easy to see that any sentence that follows from a set of Hempel-confirmed sentences is Hempel-confirmed, too.¹⁰ Hence, Hempel's confirmation criterion satisfies the Consequence Condition. The same holds true of the Consistency Condition. Indeed, Hempel's proposal satisfies his own adequacy conditions. Moreover, many intuitively clear cases of confirmation are successfully reconstructed in Hempel's account. However, one can raise several objections against Hempel, some of which were anticipated by Hempel himself in a postscript to "Studies in the Logic of Confirmation".

First, some hypotheses do not have finite developments and are therefore not confirmable. Take the hypothesis

$$H_2 = (\forall x : \neg Gxx).(\forall x : \exists y : Gxy).(\forall x, y, z : Gxy.Gyz \rightarrow Gxz)$$

which asserts that G is a serial, irreflexive and transitive two-place relation. These properties entail that H_2 is not satisfiable in any finite structure and thus not Hempel-confirmable by a finite number of observations. But certainly, H_2 is not meaningless — you might interpret G as the 'greater than' relation and then, the natural numbers with their ordinary ordinal structure are a model of H_2 . Read like this, H_2 asserts that the 'greater than' relation is transitive, irreflexive and for any natural number, there is another natural number which is greater than it. It is strange that such hypotheses are not confirmable pace Hempel. The problem is maybe purely technical, but it is nevertheless embarrassing.

Second, consider c , an individual constant of our predicate language, and the hypotheses $H_3 = \forall x : Ix$ and $H_4 = \forall x : (x \neq c \rightarrow \neg Lx)$. Take the set of all planets of the solar system as the universe of our intended structure and let the individual constant c refer to Planet Earth. Then H_3 might be interpreted as the claim that iron exists on all planets and H_4 as the claim that no life exists on other planets. Both are meaningful hypotheses open to empirical investigation. Now, the observation report $E = Ic$ (there is iron on Earth) directly Hempel-confirms $H_3.H_4$ (there is iron on all planets and life does not exist on other planets) relative to empty background knowledge.¹¹ While this may still be acceptable, it also follows that H_4 is Hempel-confirmed by $E = Ic$, due to the Special Consequence Condition. This is utterly strange since the actual observation (there is iron on Earth) is completely independent of the hypothesis at stake (no life exists on other planets). Clearly, this conclusion goes beyond what the available evidence entitles us to infer. More embarrassing, this type of inference is generalizable to other examples, too.¹²

¹⁰Assume that $\mathcal{S} \models H$ where \mathcal{S} is Hempel-confirmed by E . Then there is a set Γ so that any element of Γ is directly Hempel-confirmed by E and that $\Gamma \models \mathcal{S}$. Since by assumption $\mathcal{S} \models H$, it follows that $\Gamma \models H$, too. Thus H is Hempel-confirmed by E .

¹¹The development of $H_3.H_4$ with regard to c is Ic .

¹²Cf. [Earman and Salmon, 1992].

These technical problems may be mitigated in refined formulations of Hempel-confirmation, but there are more fundamental problems, too. They are in a similar vein connected to the fact that Hempel-confirmation satisfies the Special Consequence Condition. When a hypothesis H is Hempel-confirmed by a piece of evidence E (relative to K), any arbitrary disjunction X can be tacked to H while leaving the confirmation relation intact. For example, the hypothesis that all ravens are black *or* all doves are white is Hempel-confirmed by the observation of a black raven, although it is not clear in how far that observation is *relevant* for the hypothesis that all doves are white. Even worse, the same observation also confirms the hypothesis that all ravens are black *or no* doves are white. The tacked disjunction is completely arbitrary. Evidential relevance for the hypothesis gets lost, but a good account of confirmation should take care of these relations.

Finally, consider the following case: A single card is drawn from a standard deck. We do not know which card it is. Compare, however, the two hypothesis that the card is the ace of diamonds (H_5) and that the card is a red card (H_6). Now, the person who draws the card tells us that the card is either an ace or a king of diamonds. Obviously, the hypothesis H_6 is entailed by the evidence and thus Hempel-confirmed. But what about H_5 ? We are now much more confident that H_5 is true because the evidence *favours* the hypothesis that the card is an ace of diamonds over the hypothesis that the card is no ace of diamonds, in the usual *relative* sense of confirmation. However, the observation does not Hempel-confirm the hypothesis that the card is an ace of diamonds. This is so because not all assertions H_5 makes about this particular card — that it is an ace and a diamond — are satisfied by the observation report. This behavior of Hempel-confirmation is awkward and stands in contrast to the most popular quantitative account of confirmation, the Bayesian account. Our toy example has analogues in science, too: it is not possible to Hempel-confirm all three of Kepler's laws by confirming one of its three components. Any confirming observation report would have to entail each of Kepler's laws (with regard to the planet that is observed). This is at least strange because we often cannot check each prediction of a theory. To give an example, an observation of the diffraction pattern of light apparently confirms the hypothesis that light is an electromagnetic wave. But waves have more characteristic properties than just exhibiting a diffraction pattern — in particular, properties that are not shown in our particular observation. Partial confirmation thus becomes difficult on a Hempelian view of confirmation. Hence, Hempel's satisfaction criterion is not only liable to severe technical objections, but also fails to reconstruct an important line of thought in scientific observation and experimentation.

Thus, the above objections do not only illuminate technical shortcomings of Hempel's account, but also a general uneasiness with the Consequence Condition and the Special Consequence Condition. But why did they seem to be so plausible at first sight? I believe *pace* Carnap [1950] that the missing distinction between the *absolute* and the *relative* concept of confirmation is the culprit.¹³ We often

¹³A discussion of that criticism which is more charitable towards Hempel can be found in [Huber, 2008].

say that a certain theory is well confirmed, but we also say that a certain piece of evidence confirms a hypothesis. These two different usages correspond to different meanings of the word ‘confirmation’. When we use the former way of speaking — ‘theory T is well confirmed’ — we say something about a particular theory: T enjoys high confidence, the total available evidence speaks for T and favors it over all serious rivals. To be confirmed or to be well confirmed becomes a property of a *particular hypothesis* or theory. By contrast, the latter use says something about a *relationship* between hypothesis and evidence — it is asked whether a piece of evidence supports or undermines a hypothesis. Relative confirmation means that an empirical finding, a piece of evidence, lends support to a hypothesis or theory. This need, however, *not* imply that on account of the total available evidence, the theory is highly credible.

The Consequence Condition is plausible whenever absolute confirmation is examined. When a strong, comprehensive theory is strongly endorsed — in the sense of ‘highly plausible’ or ‘empirically supported beyond all reasonable doubt’ — any part of this theory is also highly plausible, etc., in agreement with (CC) and (SCC). Obviously, the less risky a conjecture is, the more confidence can we put in it, and any proper part of a theory is logically weaker and thus less risky than the entire theory. Therefore the Consequence Condition makes perfect sense for degrees of belief and conviction, i.e. when it comes to endorsement and absolute confirmation. It is, however, highly questionable whether the Consequence Condition is also a sensible condition with regard to relative confirmation. Here, the evidence has to be *informative* with respect to the hypothesis under test. For instance, Eddington’s observations of the 1919 eclipse apparently confirmed the hypothesis that light is bent by massive bodies as the sun. General Theory of Relativity (GTR), the overarching theory, was at that time still fiercely contested, and the agreement of Eddington’s observations with GTR and their discrepancy from the Newtonian predictions constituted key evidence in favor of GTR. (The bending effect in the GTR predictions was roughly twice as high as in Newtonian theory.) But it would be much more controversial to claim — as (CC) does — that Eddington’s observations directly confirmed those parts of GTR that were remote from the bending-of-light effect, e.g. the gravitational redshift which was proven in the Pound-Rebka experiment in 1959. Confirmation does not automatically transmit to other sub-parts of an overarching theory, as vindicated in the (probabilistic) analysis by Dietrich and Moretti [2005]. Thus, we are well advised to drop the Consequence Condition. A similar criticism can be directed against the Consistency Condition since *any* coherent and unified theory that were in agreement with Eddington’s observations would have been confirmed by them.

In a postscript to “Studies in the Logic of Confirmation” that appeared in 1965, Hempel admitted some of the problems of his account. In particular, he felt uncomfortable about the Consistency Condition which he thought to be too strong to figure as a necessary condition for (relative) confirmation. Thus, the satisfaction criterion is too narrow as a qualitative definition of confirmation. This concession suggests that Hempel actually spotted the problem of combining the concepts of

relative and absolute confirmation in a single account (cf. [Huber, 2008]). But Hempel [1945/1965, p. 50] still contends that his adequacy conditions may be *sufficient* for a definition of confirmation. However, the next section will come up with a telling counterexample — a case of spurious confirmation which the satisfaction criterion fails to discern.

4 THE RAVEN PARADOX

Hypotheses about natural laws and natural kinds are often formulated in the form of universal conditionals. For instance, the assertion that all F 's are G 's ($H = \forall x : Fx \rightarrow Gx$) suits hypotheses like 'all planets have elliptical orbits', 'all ravens are black' or 'all cats are predators'. How are such claims confirmed? There is a longstanding tradition in philosophy of science that stresses the importance of *instances* in the confirmation of universal conditionals, going from Nicod [1925/1961] over Hempel [1945/1965] to Glymour [1980]. A confirming instance consists in the observation of an F that is also a G ($Fa.Ga$) whereas an observation of an F that is no G ($Fa.\neg Ga$) refutes H . According to Nicod, only these two kinds of observation — 'confirmation' and 'infirmation' — are relevant to the hypothesis. *L'induction par l'infirmation* proceeds by refuting and eliminating other candidate hypothesis, *l'induction par la confirmation* supports a hypothesis by finding their instances. There is, however, an important asymmetry [Nicod. 1925/1961, pp. 23-25]: while observing a non-black raven refutes the raven hypothesis once and for all, observing a black raven does not permit such a conclusive inference. Nicod adds that not the sheer number of instances is decisive, but the *variety* of instances which can be accrued in favor of that hypothesis.

If we try to put this idea of instance confirmation into a single condition, we might arrive at the following condition:

Nicod Condition (NC): For a hypothesis of the form $H = \forall x : Rx \rightarrow Bx$ and an individual constant a , an observation report of the form $Ra.Ba$ confirms H .

However, this account does not seem to exhaust the ways a hypothesis can be confirmed. Recall the

Equivalence Condition (EC): If H and H' are logically equivalent sentences then E confirms H if and only if E confirms H' .

As already argued, the equivalence condition is an uncontroversial constraint on a logic of confirmation. Combining (EC) with Nicod's condition about instance confirmation leads, however, to paradoxical results: Take the hypothesis that nothing that is non-black can be a raven ($H' = \forall x : \neg Bx \rightarrow \neg Rx$). A white shoe is an instance of that hypothesis, thus, observing it counts as a confirming observation report. By the Equivalence Condition, H' is equivalent to $H = \forall x : Rx \rightarrow Bx$ so that a white shoe also confirms the hypothesis that all ravens are black.

But a white shoe seems to be utterly irrelevant to the color of ravens. Hence, we have three individually plausible, but incompatible claims at least one of which has to be rejected:

1. **Nicod Condition (NC):** For a hypothesis of the form $H = \forall x : Rx \rightarrow Bx$ and any individual constant a , an observation report of the form $Ra.Ba$ confirms H .
2. **Equivalence Condition (EC):** If H and H' are logically equivalent sentences then E confirms H relative to K if and only if E confirms H' relative to K .
3. **Confirmation Intuition (CI):** A Hypothesis of the form $H = \forall x : Rx \rightarrow Bx$ is not confirmed by an observation report of the form $\neg Ra.\neg Ba$.

This set of jointly inconsistent claims constitutes the *paradox of confirmation* and was first discussed in detail by Hempel [1965].¹⁴ The main conflict consists in the fact that (EC) and (NC) merely consider the *logical form* of scientific hypotheses whereas (CI) implicitly assumes that there is an ‘intended domain’ of a scientific hypothesis. In particular, only ravens seem to be evidentially relevant to the hypothesis that all ravens are black.

One option to dissolve the paradox discussed (and rejected) by Hempel [1945/1965] consists in re-interpreting the hypothesis. General natural laws in the form of universal conditionals apparently confer existential import on the tentative hypotheses: ‘All ravens are black’ could be read as ‘all ravens are black and there exists at least one raven’. Then, there is no inconsistency between the above three claims. But that proposal is not convincing. The observation of a single black raven provides conclusive evidence in favor of the second part of the hypothesis. As Alexander [1958, p. 230] has pointed out, we will then focus on confirming or undermining the first part of the hypothesis (‘all ravens are black’) as soon as a black raven has been observed. Hence, the paradox appears again. Interpreting the raven hypothesis as having existential import does not remove the problem.

Before going into the details of attempted solutions it is interesting to note a line of thought that can be traced back to Hempel himself.

“If the given evidence E [...] is black [Ba], then E may reasonably be said to even confirm the hypothesis that all objects are black [$\forall x : Bx$], and *a fortiori*, E supports the weaker assertion that all ravens are black [$H = \forall x : Rx \rightarrow Bx$].”¹⁵

We can transfer this argument in a canonical to non-ravens (cf. [Goodman, 1983, pp. 70-71]):

¹⁴Note that the inconsistency vanishes if the conditionals are interpreted as *subjunctive* and not as material conditionals: contraposition is not a valid form of inference for subjunctive conditionals.

¹⁵[Hempel, 1945/1965, p. 20].

If the given evidence E is a non-raven $[\neg Ra]$, then E may reasonably be said to even confirm that all objects are non-ravens $[\forall x : \neg Rx]$, and *a fortiori*, E supports the weaker assertion that all non-black objects are non-ravens $[\forall x : \neg Bx \rightarrow \neg Rx]$, *i.e.* that all ravens are black $[H = \forall x : Rx \rightarrow Bx]$.¹⁶

Thus we obtain another ostensibly decisive argument against (CI). But as remarked by Fitelson [2006], the argument requires additional assumptions. Above all, the step from ‘a non-raven confirms H ’ to ‘a *black* non-raven confirms H ’ is far from trivial — it rests on the principle of *monotonicity* that extending the evidence cannot destroy the confirmation relation. Without this additional claim, the above argument would not bear on the observation of a black non-raven. Moreover, Hempel’s adequacy condition (SCC) is employed, namely in the transition from ‘ E confirms $\forall x : \neg Rx$ ’ to ‘ E confirms $\forall x : \neg Bx \rightarrow \neg Rx$ ’. We may suspend judgment on monotonicity, but (SCC) is, as seen in the previous section, a controversial condition on relative confirmation. So the above reasoning does not remove the paradox convincingly.¹⁷

Hempel suggests that we should learn to live with the paradoxical conclusion. His argument can be paraphrased thus:¹⁸ Assume that we observe a grey, formerly unknown bird that is in most relevant external aspects very similar to a raven. That observation puts the raven hypothesis to jeopardy. It might just be the case that we have seen a non-black raven and falsified our hypothesis. But a complex genetic analysis reveals that the bird is no raven. Indeed, it is more related to crows than to ravens. Hence, it sounds logical to say that the results of the genetic analysis corroborate the raven hypothesis — it was at risk and it has survived a possible falsification. In other words, a potential counterexample has been eliminated. Thus there is no paradox in saying that an observation report of the form $\neg Ra. \neg Ba$ confirms H , in the sense that a satisfies the constraint given by H that nothing can be both a raven and have a color different from black.¹⁹

Hempel elaborates the crucial point in more detail, too: Compare two possible observation reports. First, we observe a crow *which we know to be a crow* and notice that it is grey ($E_1 = \neg Ba$, $K_1 = \neg Ra$). This seems to be a fake experiment

¹⁶I borrow the idea to paraphrase Hempel’s argument in this way from Maher [1999] and Fitelson [2006].

¹⁷Quine [1969], by contrast, defends (CI) and finds the paradox unacceptable. Since he maintains (EC), too, he is forced to reject the Nicod Condition. Nonetheless, he defends a modified Nicod Condition whose content is restricted to *natural kinds*. Only instances of natural kinds confirm universal conditionals, and clearly, neither non-ravens nor non-black things count as natural kinds. However, this line of reasoning is subject to the Hempelian criticism explained in the text.

¹⁸[Hempel, 1945/1965] makes the argument for quite a different example (‘all sodium salts burn yellow’) but I would like to stick to the original raven example in order not to confuse the reader.

¹⁹It might now be objected that the observation of a black raven seems to lend *stronger* support to the raven hypothesis than the observation of a grey crow-like bird since such an observation is more relevant to the raven hypothesis. But this is a problem for a quantitative account of confirmation (we will get back to this in section 5) and not for a qualitative one.

if evaluated with regard to the raven hypothesis — we knew beforehand that a crow could not have been a non-black raven. There was no risk involved in the experimentation, so neither confirmation nor Popperian corroboration could result. In the second case we observe an object *about which we do not know anything beforehand* and discover that the bird is a grey crow ($E_2 = \neg Ra. \neg Ba$, $K_2 = \emptyset$). That counts as a sound case of confirmation, as argued above. Hempel describes the difference thus: When we are told beforehand that the bird is a crow

[...] “this has the consequence that the outcome of the [...] color test becomes entirely irrelevant for the confirmation of the hypothesis and thus can yield no new evidence for us.”²⁰

In other words, the available background knowledge in the two cases makes a crucial difference. Neglecting this difference is responsible for the fallacious belief (CI) that non-black non-ravens cannot confirm the hypothesis that all ravens are black. (CI) is plausible only if we tacitly introduce the additional *background knowledge* that the test object is no raven. Thus, in the above example, H should be confirmed if we do *not* know beforehand that the bird under scrutiny is a crow ($K_2 = \emptyset$) and it should not be confirmed if we *know* beforehand that the bird is a crow ($K_1 = \neg Ra$). In Hempel’s own words,

“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) [...] we have to ask: Given some object a [that is neither a raven nor black, but we do not happen to know this, J.S.]: does a constitute confirming evidence for the hypothesis? And now [...] it is clear that the answer has to be in the affirmative, and the paradoxes vanish.”²¹

Thus, the paradox is a psychological illusion, created by tacit introduction of background knowledge into the confirmation relation. From a logical point of view, (CI) reveals itself as plainly false. One of the three premises of the paradox has been discarded. A problem that remains, though, is that the Hempelian resolution does not make clear why ornithologists should go into the forest to check their hypothesis and not randomly note the properties of whatever object they encounter. This might be called the problem of *armchair ornithology*.

In fact, this criticism is raised by Watkins [1957]. Watkins insinuates that Hempel may cheaply confirm $H = \forall x : Rx \rightarrow Bx$ by summing up observations of non-black non-ravens while sitting in the armchair.²² In a similar vein, Watkins [1957] objects that cases of confirmation as the observation of a white shoe do not

²⁰[Hempel, 1945/1965, p. 19].

²¹[Hempel, 1945/1965, pp. 19-20].

²²A reply to Watkins is [Vincent, 1964].

put the hypothesis to a real test and thus contradict the falsificationist methodology for scientific hypotheses. On the Popperian, falsificationist account, hypotheses can only be corroborated by the survival of severe tests, and observing shoes does not count as a real test of a hypothesis. Second, the ‘negation’ of observing a white shoe, namely observing a black shoe, would equally confirm the raven hypothesis on Hempel’s account. This trivializes the notion of instance confirmation on which Hempel’s satisfaction criterion is based. Every universal conditional is automatically confirmed by lots of irrelevant evidence. Watkins concludes that an inductivist reasoning about confirmation (as Hempel’s instance confirmation) be better replaced by a truly falsificationist account.

Alexander [1958; 1959] answers to Watkins that falsificationist corroboration also presupposes some kind of inductive reasoning if it is supposed to affect our expectations on the future: If a hypothesis survives several tests, we expect that it will survive future tests, too — otherwise it would not make sense to say that the hypothesis has been corroborated. So Watkins’s dismissal of inductive, instance-based reasoning goes too far. Moreover, Hempel makes an important proviso, namely that there be no substantial background assumptions when evaluating the evidential relevance of $\neg Ra. \neg Ba$. If we do not know where to find and where not to find ravens, i.e. if we randomly sample from the class of *all* objects, the observation of a white shoe *does* count as a genuine test of the raven hypothesis. One may object that Hempel’s proviso is unrealistic for actual cases of (dis)confirmation (cf. [Watkins, 1960]), but conditional on this proviso, Hempel’s conclusion — everything that is not a non-black raven supports $H = \forall x : Rx \rightarrow Bx$ — seems to be correct. So the first objection vanishes. Second, it is misleading to say that the raven hypothesis is confirmed by conflicting evidence — rather, *different kinds of evidence* (namely, shoes of different color) equally confirm the hypothesis. Similarly, observing male as well as female black ravens confirms the raven hypothesis. Here, nobody would object that those pieces of evidence are conflicting and therefore inadmissible for confirmation.

However, as pointed out by Agassi [195]), Hempel’s conclusion is stronger than that a non-black non-raven *may* confirm the raven hypothesis — it is claimed that this piece of evidence *always* confirms the raven hypothesis, independent of the background knowledge. Good [1960; 1961] has suggested the following (slightly modified) example to refute that conjecture: The only black middle-sized objects which a child sees are black crows and black ravens. No other black objects occur, and all ravens and crow the child sees are black. Suddenly she discovers a white crow. Then she says: “How surprising! Apparently objects that are supposed to be black can sometimes be white instead.”²³ And what is good for the goose (crows) is equally good for the gander (ravens). So the child concludes that ravens may be white, too. On Hempel’s account, the observation of a grey crow would support rather than undermine the hypothesis that all ravens are black. Isn’t that behavior insensitive to the peculiarities of the specific case?

I believe Agassi and Good are on the right track, but they do not fully pin down

²³[Good, 1961, p. 64]. Cf. [Swinburne, 1971].

Hempel's problem. We may admit that Hempel succeeds in explaining away the paradoxical nature of the problem. But his own satisfaction criterion fails to resolve the paradox. Remember Hempel's diagnosis that tacitly introduced or deliberately suppressed background information is the source of the paradox. While perfectly agreeing with Hempel on this point, Fitelson and Hawthorne [2009] point out that Hempel is unable to make that difference in his own theory of confirmation. The reason is that his account is in general *monotone* with regard to the background knowledge: As long as the domain of the evidence is not extended (i.e. no individual constants are added), additional background knowledge cannot destroy the confirmation relation. Hempel inherits this property from deductive logic, because $E.K \models H_{|dom(E)}$ is the crucial condition for direct Hempel-confirmation, and thus also for Hempel-confirmation. Evidently, logical entailment is preserved under adding additional conditions to the antecedens. Therefore Hempel's own account yields confirmation even if the background knowledge is far too strong. In the first case (we do not know beforehand that a is no raven) confirmation follows from

$$E_1.K_1 = \neg Ra.\neg Ba \models (Ra \rightarrow Ba) = H_{|dom(E)}$$

and in the second case, we have precisely the same implication

$$E_2.K_2 = \neg Ra.\neg Ba \models (Ra \rightarrow Ba) = H_{|dom(E)}.$$

Hence, adding the background knowledge that the test object is no raven does not destroy the (Hempel-)confirmation of H_2 . Certainly Hempel spots two points correctly: First, the paradoxical conclusion of the raven example should be embraced, contra (CI). Second, background knowledge plays a crucial role when it comes to explaining the source of the paradox. But while pointing into the right direction, Hempel fails to set up an account of confirmation that conforms to his own diagnosis of the paradox. In particular, the adequacy criteria outlined in section 2 fail to be sufficient for a satisfactory concept of confirmation.

The raven paradox drastically shows how valuable it is to distinguish between evidence and background knowledge. The distinction has to be formalized in a way that avoids Hempel's problem. It further exhibits the problem of monotonicity with regard to evidence and background knowledge: When we happen to know more, confirmation might get lost. Therefore monotonicity is not a desirable property for accounts of confirmation, and I take this to be the third important moral from the paradoxes of confirmation. On the other hand, the arguments to resolve the paradox by giving up (CI) were on a whole convincing, and Hempel's sixty-three-year-old judgment that part the paradoxical appearance often rests on a psychological illusion has some plausibility. The next section examines the paradoxes of confirmation from a probabilistic perspective.²⁴

²⁴In a recent paper, Branden Fitelson [2009] elaborates the similarity of the raven paradox to a famous logical puzzle: the Wason Selection Task [Wason and Shapiro, 1971]. In the Wason Selection Task, four cards lie on the table. On the front side of each card, there is a letter, on the back side, there is a number. The hypothesis H is: All cards with an even number on

5 THE BAYESIAN'S RAVEN PARADOX

5.1 Bayesian confirmation theory and the Nicod Condition

So far, we have discussed the paradox in a qualitative way — does observing a non-black non-raven confirm the hypothesis that all ravens is black? The Hempelian resolution does, however, not clarify why we would recommend an ornithologist to go into the forest, in order to confirm the raven hypothesis. A natural reply would contend that black ravens confirm the raven hypothesis to a much stronger degree than white shoes. That thesis motivates a quantitative treatment of the paradox and will be the main subject of this section. Actually, the ‘confirmation intuition’ (CI) about the missing confirmatory value of non-ravens has three versions — a qualitative, a comparative and a quantitative one:

Qualitative Intuition The observation of a non-black non-raven does not confirm the hypothesis that all ravens are black.

Comparative Intuition The observation of a non-black non-raven confirms the hypothesis that all ravens are black to a lower degree than the observation of a black raven.

Quantitative Intuition The observation of a non-black non-raven confirms the hypothesis that all ravens are black only to a minute degree.

Part of the confusion in the existing literature is due to the fact that these three intuitions are not clearly set apart from each other. Hempel criticized exclusively the qualitative version. The quantitative and the comparative versions save the part of (CI) that concerns the *extent* of confirmation, and here our intuitions seem to be more stable. They form the resilient kernel of (CI) which makes the raven paradox so intriguing for modern confirmation theory.

A further source of confusion is the question which background knowledge should be taken when evaluating these intuitions. Are they meant to hold for *some*, for *empty* or for *all* conceivable background assumptions? Or are those intuitions relative to the *actual* background assumptions?²⁵ Hence, twelve ($=3 \times 4$) different confirmation intuitions about the paradox could in principle be distinguished. But I believe intuitions with respect to *actual* background knowledge to be most interesting. First, most people seem to have that in mind when being

one side have a vowel printed on the other side. Which of the cards ($A, 2, F, 7$) should you turn over to test the truth of H ? Of course you have to turn over the card with the ‘2’ since this can be an obvious instance or counterexample to H . This line of reasoning is captured in the Nicod Condition, too. It is less obvious that you also have to turn over the ‘ F ’ in order to test the contrapositive: All cards with a consonant on one side have an odd number on the other side. People regularly fail to recognize that the ‘ F ’ has to be turned over, too. The kind of confirmation which this action yields is structurally identical to confirming the raven hypothesis by observing that a grey bird is not a raven, but a crow. Both the results in the Wason Selection Task and the debate around the raven paradox highlight the same kind of reluctance to accept instances of the contrapositive as instances of the hypothesis itself.

²⁵I borrow these distinctions from [Fitelson, 2006].

confronted with the paradox, so it is arguably the most accurate reconstruction of the paradox. Second, we will later argue that the Nicod Condition is best understood as referring to actual background knowledge. Indeed, Good's [1961] raven/crow example suggests that the above confirmation intuitions will trivially hold for *some* background knowledge and trivially be false for *every* conceivable background knowledge. Finally, what *empty* background knowledge means stands in need of explication (though see [Carnap, 1950; Maher, 2004]). Thus we are well advised to focus on *actual* background knowledge. Here we have seen that the qualitative version of (CI) is under pressure, but on the other hand, the comparative and the quantitative versions enjoy some plausibility. This section tries to reinforce the arguments against the qualitative intuition and to vindicate the comparative and quantitative intuition from the point of view of Bayesian confirmation theory. The problem with the raven paradox is not the alleged truth of (CI), but the truth of the weaker comparative and quantitative versions.

Qualitatively, Bayesian confirmation amounts to an increase in rational degree of belief upon learning new evidence. Degrees of belief are symbolized by subjective probabilities. In other words, evidence E confirms H if and only if $P(H|E) > P(H)$. But we have to remember a lesson from the very first chapter of the book — confirmation is a three place predicate, relative to background knowledge. As both the raven paradox and the Duhem problem teach us, background assumptions are a crucial part of relating theory to evidence and inductive reasoning in science. The natural way to integrate them consists in taking background information for granted and conditionalizing an agent's degrees of belief on it.²⁶ That said, we can write down a first, qualitative definition of Bayesian confirmation:

DEFINITION 9. A piece of evidence E *confirms* a hypothesis H relative to background assumptions K if and only if $P(H|E.K) > P(H|K)$.

This definition gives a probabilistic explication of *relative confirmation*, not of *absolute confirmation*: Definition 9 describes the relevance of evidence for a hypothesis, not high credibility of a hypothesis. However, the definition remains qualitative. To be able to tackle the comparative and quantitative versions of the paradox, we have to introduce a *measure of confirmation*. The following three candidates have been especially popular in the literature (see Fitelson 2001 for a discussion of their virtues and vices):

Difference Measure

$$d(H, E, K) := P(H|E.K) - P(H|K)$$

Log-Ratio Measure

$$r(H, E, K) := \log \frac{P(H|E.K)}{P(H|K)}$$

²⁶Nonetheless, for reasons of convenience, we will often speak (but not write) as if the background knowledge were empty.

Log-Likelihood Measure

$$l(H, E, K) := \log \frac{P(E|H.K)}{P(E|\neg H.K)}$$

For reasons of simplicity, I restrict myself in the following to d and l which suffice to illustrate the substantial points.

In the 1950s and early 1960s, the discussion of the confirmation paradoxes focussed on discussing, defending and rebutting (CI). In particular, Hempel himself has rejected (CI) and argued that tacit introduction of background knowledge may be responsible for the paradoxical appearance. In the light of Bayesian confirmation theory, one could, however, not only reject (CI), but also question (NC). Again, four versions of (NC) have to be distinguished.

Nicod Condition (NC): For a hypothesis of the form $H = \forall x : Rx \rightarrow Bx$ and any individual constant a , an observation report of the form $Ra.Ba$ confirms H , relative to every/actual/tautological/any background knowledge.

Certainly, the Nicod Condition (every black raven confirms the raven hypothesis) is true relative to *some* background knowledge. But that claim is very weak and practically not helpful. It is somewhat more surprising that it is not true under *all* circumstances. I. J. Good [1967] constructed a simple counterexample in a note for the *British Journal for the Philosophy of Science*: There are only two possible worlds. In one of them, W_1 , there are a hundred black ravens, no non-black ravens and one million other birds. In the other world W_2 , there are a thousand black ravens, one white raven and one million other birds. Thus, H is true whenever W_1 is the case, and false whenever W_2 is the case. For all suggested measures of confirmation, the observation of a black raven is evidence that W_2 is case and therefore evidence that not all ravens are black:

$$P(Ra.Ba|W_1) < P(Ra.Ba|W_2)$$

and consequently,

$$l(H, E, K) = \log \frac{P(E|W_1)}{P(E|W_2)} \approx \frac{100}{10^6} \frac{10^6}{1000} < 0 \quad (1)$$

Hence, (1) refutes the claim that all universal conditionals are confirmed by their positive instances. Although the example is highly artificial, it refutes the strong version of (NC) that is generalized over all conceivable circumstances. In his reply to Good, Hempel [1967] explained that he intended the evidential relevance relation in the raven paradox (and thus also the Nicod Condition) to be “considered by itself and without reference to any other information”.²⁷ So he opts for the version of (NC) relative to empty, or tautological, background knowledge. Let’s

²⁷[Hempel, 1967, pp. 239-40].

call this claim (NC_T) . Good [1968] objects, however, in the best Humean tradition, that it is hard to make sense of (NC_T) : In the absence of *any* empirical background knowledge, the raven hypothesis itself and the pieces of evidence are apparently meaningless. He concludes that in such circumstances, it is hard (or even impossible) to decide whether a black raven confirms the hypothesis that all ravens are black. In a more nostalgic mood, one could contend that Hempel, by his adherence to (NC_T) , re-introduced the *a priori* into empiricist philosophy of science: It must, or so he has to argue, be clear from the meaning of ‘black’ and ‘raven’ that, in the absence of further information, observing a black raven confirms the hypothesis that all ravens are black. But even if we play Hempel’s game and suppress empiricist concerns, (NC_T) remains questionable. In his [2004], Patrick Maher has explicated (NC_T) in a Carnapian framework of inductive probability and given a counterexample, i.e. he has shown that relative to a certain way of framing tautologous background knowledge, the observation of a black raven does *not* support the claim that all ravens are black. One may, of course, contest Maher’s explication of tautological background knowledge, e.g. the choice of parameters in his logical inductive framework. But then one has to defend a better explication, too, and so far, this has not been done.

In order not to trivialize the paradox of confirmation, it should thus be understood relative to our *actual* background knowledge. In particular, (NC_A) , the ‘actual’ version of the Nicod condition, seems to be the most sensible interpretation of (NC) . No one will seriously doubt that observing a black raven actually confirms the raven hypothesis.²⁸ Now we return to the original questions, namely: What is the status of the comparative and quantitative versions of (CI) ?

5.2 The quantitative question

Let’s begin with the quantitative question.²⁹ A number of approaches (e.g. [Alexander, 1959; Mackie, 1963]) start with the assumption that the ratio of ravens versus non-ravens and black things versus non-black things is fixed, regardless of whether all ravens are black. On the behalf of this and another independence assumptions, they are able to show that black ravens confirm the raven hypothesis better than white shoes, as measured by any standard Bayesian measure of confirmation. But the initial assumption is very strong and also not very plausible — why shouldn’t the discovery of a population of white ravens increase our estimate of the total number of ravens? Even worse, those authors admit that black non-ravens *disconfirm* H although black non-ravens also conform to H . There is no reason at all why a white shoe — or more generally, a white non-raven — should be positively relevant to H whereas a black shoe is negatively relevant to H . To my mind,

²⁸Although the truth of (NC_A) seems to be obvious, it is in fact difficult to *prove*: there are no canonical probability distributions for our actual situation. This problem transfers, as we will see, to the probabilistic analysis of the paradoxes of confirmation in general.

²⁹There are a lot of related quantitative approaches that cannot be covered in this article, including [Black, 1966; Gaifman, 1979; Hossiasson-Lindenbaum, 1940; Humburg, 1986; Maher, 1999; Royall, 1997; von Wright, 1966].

such a ‘resolution’ is nothing more than a *reductio ad absurdum* of the initial assumptions.

Other approaches try to dissolve the paradox by representing the evidence-gathering process in two stages (e.g. [Horwich, 1982]). For instance, we apparently obtain more confirmation from observing that a (known) raven is black than from observing that a (known) white object is a shoe. This is an interesting way to tackle the paradox — it actually facilitates a Bayesian resolution —, but it alters the question which was concerned with the *total* amount of confirmation yielded by $Ra.Ba$ and $\neg Ra.\neg Ba$. Hence, those approaches do not directly address the quantitative confirmation intuition.

Howson and Urbach [1993] argue with Mackie [1963] that black ravens are less frequently found than non-black non-ravens (e.g. white shoes), and that for this reason, black ravens are more *informative* evidence than non-black non-ravens.

“[...] since non-black non-ravens form such a numerous class compared with black ravens, it is almost [...] certain that a random object about which we know nothing will turn out to be neither black nor a raven, but relatively unlikely that it will be a black raven.”³⁰

They conclude that black ravens provide much better confirmation than non-black non-ravens. But that is a mistake, as pointed out by Korb [1994]. The probability of the evidence *by itself* (i.e. $P(Ra.Ba)$ vs. $P(\neg Ra.\neg Ba)$) is not sufficient to settle the degree of confirmation. Rather, what we have to compare are the probabilities of the evidence *under the competing hypotheses*, i.e. $P(Ra.Ba|H)$, $P(Ra.Ba|\neg H)$, etc. So Howson and Urbach’s argument does not get off the ground.

Let’s now get to the standard Bayesian approach to the paradox. It requires two assumptions:

1. *Proportion of ravens assumption (PRA)*: Among the non-black objects, there are almost no ravens, i.e. $P(Ra | \neg Ba.K) \approx 0$.³¹ — We justify (PRA) by noting that the non-black objects in the world clearly outnumber the ravens, and in particular the non-black ravens. Our background knowledge ensures that we will not discover large non-black raven populations. Even in worlds where H is false we will not encounter many non-black ravens, whereas there will be a huge amount of non-black objects in any case.
2. *Independence assumption (IA)*: The color of an individual a does not depend on the truth of the raven hypothesis, i.e. $P(Ba | H.K) \approx P(Ba|K)$. — (IA) is harder to defend than (PRA). For instance, if a is a raven, the falsity of H seems to raise the the likelihood of a being black. (IA) claims that nevertheless, the overwhelming number of non-black objects and the comparably small number of ravens ensure that such effects are marginal.

³⁰[Howson and Urbach, 1993, p. 90].

³¹Note that this proposition also makes a claim about worlds in which the raven hypothesis is *false*.

Let's see how far we get if we assume (IA) and (PRA). Let $\varepsilon = P(Ra \mid \neg Ba.K)$, $\delta = P(\neg Ba \mid K) - P(\neg Ba \mid H.K)$. Now we try to measure the amount of confirmation which $\neg Ra.\neg Ba$ adds to the raven hypothesis. By Bayes' Theorem and the *multiplication rule* $P(A.B \mid C) = P(A \mid B.C)P(B \mid C)$ we obtain

$$\begin{aligned} & d(H, \neg Ra.\neg Ba, K) \\ \stackrel{df}{=} & P(H \mid \neg Ra.\neg Ba.K) - P(H \mid K) \\ = & P(H \mid K) \left[\frac{P(\neg Ra.\neg Ba \mid H.K)}{P(\neg Ra.\neg Ba \mid K)} - 1 \right] \\ = & \frac{P(H \mid K)}{P(\neg Ra.\neg Ba \mid K)} [P(\neg Ra \mid \neg Ba.H.K)P(\neg Ba \mid H.K) - P(\neg Ra.\neg Ba \mid K)] \\ = & \frac{P(H \mid K)}{P(\neg Ra.\neg Ba \mid K)} [P(\neg Ba \mid H.K) - P(\neg Ra \mid \neg Ba.K)P(\neg Ba \mid K)] \end{aligned}$$

Then we use the definitions of ε and δ and obtain

$$\begin{aligned} & d(H, \neg Ra.\neg Ba, K) \\ = & \frac{P(H \mid K)}{P(\neg Ra.\neg Ba \mid K)} [P(\neg Ba \mid K) - \delta - (1 - \varepsilon)P(\neg Ba \mid K)] \\ = & \frac{P(H \mid K)}{P(\neg Ra.\neg Ba \mid K)} [\varepsilon P(\neg Ba \mid K) - \delta] \\ = & O(\varepsilon, \delta) \end{aligned}$$

(IA) and (PRA) control ε and δ . Hence, the calculation shows that $d(H, \neg Ra.\neg Ba, K) \approx 0$ for sufficiently small ε and δ , solving the quantitative puzzle. Indeed, our intuitions tell us that the confirming effect of eliminating possible counterexamples has to be marginal. Moreover, the calculation explains why we first thought that such observations cannot deliver any confirmation at all: we confounded marginal and null confirmation.

Although the treatment of the raven paradox argument is often judged as a success story of Bayesianism (cf. [Earman, 1992]), it is open to various objections, and the most crucial of them are directed against the premise (IA), cf. [Vranas, 2004]. After rebutting arguments in favor of (IA) given by Horwich [1982] and Woodward [1985], Vranas points out that most of the arguments use a kind of 'negative approach': we have no positive arguments that $P(Ba \mid H.K)$ be different from $P(Ba \mid K)$, so we conclude that they are approximately equal. Vranas believes the latter step to be invalid (and I side with him here):

"it does not follow that my estimate [of black objects] should remain the same [when learning H or $\neg H$]. What follows instead is that [my]

estimate *may* remain the same. [...] So even if there is no reason why $P(Ba|H.K)$ and $[P(Ba|K), J.S.]$ should differ, maybe there is no reason why they should be equal either.”³²

Indeed, given the various ways H can be true or false, it is hard to justify that in a world with grey, white or red ravens the proportion of black objects remains unchanged. These effects are the more pronounced the more narrow we conceive our reference class (i.e. our ‘universe’ of objects). Hence, our actual knowledge does not rule out that (IA) be violated.

Vranas further argues that (IA) seems to be “for all practical purposes *necessary* for the Bayesian claim [that the quantitative paradox can be resolved].”³³. Indeed, the quantitative approaches of which I am aware rely on (IA) or even more contentious claims. So we get a pessimistic result — the degree of confirmation $\neg Ra.\neg Ba$ lends to H need not be close to zero and the quantitative problem remains hard to solve. Korb [1994, p. 145] also suggests a reason for the continued failures: for all measures of confirmation, we have to incorporate, directly or indirectly, the probability of the evidence under the competing hypotheses. This requires us to assume a specific set of alternative hypotheses to ‘all ravens are black’, such as ‘most ravens are black’, ‘only some ravens are black’, etc. In particular, those alternatives have to clearly assert how they fix the proportions of ravens and black objects in an alternative scenario. The problem for all general solutions of the paradox consists in the requirement that they be compatible with all ways to carve and to weigh the space of alternative hypotheses. Furthermore, the proper reference class for the probabilistic analysis is rarely specified, i.e. it is not clear whether the calculations are supposed to hold in a universe of middle-sized birds, birds in general, or all objects in general. Hence, we need strong assumptions as (IA) if the resolution is to hold in the desired generality. To my mind, this is the reason why we have so much trouble with providing a general resolution that does not rely on contentious ad hoc presuppositions.

5.3 The comparative paradox

Given the problems with the quantitative version of the paradox, resolving the comparative solution is maybe the best we can achieve. Of course, the comparative version remains silent about the *extent* to which black ravens are more relevant than white shoes. But comparative results can at least substantiate some of our intuitions about the paradox. The most recent and ambitious result in the area is due to Fitelson and Hawthorne [2009]. They have argued that with the help of two apparently innocent assumptions, the raven paradox can be solved:

THEOREM 10. (Fitelson and Hawthorne) Let K denote the background assumptions, $H = \forall x : Rx \rightarrow Bx$. Suppose that

³²[Vranas, 2004, p. 550].

³³[Vranas, 2004, p. 548]. Italics in the original.

- $0 < P(H|Ba.Ra.K) < 1$, $0 < P(H|\neg Ba.\neg Ra.K) < 1$, $P(\neg Ba.\neg Ra|K) > 0$, $P(\neg Ba.Ra|K) > 0$, and $P(Ba.Ra|K) > 0$.

- $P(\neg Ba|H.K) > P(Ra|H.K)$.

-

$$\left(\frac{P(H | Ra.K)}{1 - P(H | Ra.K)} \right) \cdot \left(\frac{1 - P(H | \neg Ba.K)}{P(H | \neg Ba.K)} \right) \\ > P(Ba | Ra.\neg H.K) + [1 - P(Ba | Ra.\neg H.K)] \frac{P(Ra | H.K)}{P(\neg Ba | H.K)} \quad (2)$$

Then $l(H, Ba.Ra, K) > l(H, \neg Ba.\neg Ra, K)$, i.e.

$$(3) \quad \log \frac{P(Ba.Ra|H.K)}{P(Ba.Ra|\neg H.K)} > \log \frac{P(\neg Ba.\neg Ra|H.K)}{P(\neg Ba.\neg Ra|\neg H.K)}$$

and in particular

$$(4) \quad P(H | Ba.Ra.K) > P(H | \neg Ba.\neg Ra.K)$$

(The proof can be found in [Fitelson and Hawthorne, 2009].) The theorem asserts that the degree of support which $Ba.Ra$ lends to H , as measured by the log-likelihood ratio l , exceeds the degree of support $\neg Ba.\neg Ra$ lends to H (see (3)). In other words, Fitelson and Hawthorne vindicate the comparative version of (CI): black ravens confirm the raven hypothesis better than white shoes. It follows easily that the posterior probability of H is higher if a black raven is observed than if a white shoe (or any non-black raven) is observed.

To evaluate their result, we have to look at the assumptions of the theorem. The first set of assumptions is fully unproblematic: It is demanded that neither the observation of a black raven nor the observation of a non-black non-raven will determine the truth or falsity of H . Moreover, the rational degree of belief that a non-black ravens, black ravens or non-black non-ravens will be observed has to be higher than zero (though it can be infinitely small). These are just assumptions that reflect the openness of our probability assignments to empirical evidence. The second assumption is a little bit richer in content, but still extremely plausible: If H is true then we are more likely to observe a non-black object than a raven. That reflects the belief that there are many non-black objects (grey birds, for example), but comparably few ravens.

Thus the last inequality (2) carries the main burden of the theorem. Is it a plausible assumption? Let us have a look at the right hand side first. Even if H is wrong, we expect the number of black ravens to vastly exceed the number of non-black ravens. (Note that we have already observed many black ravens!) Thus, $x := P(Ba | Ra.\neg H.K)$ is quite close to 1. Moreover, in any case there are many more non-black things than ravens. So the ratio $P(Ra | H.K)/P(\neg Ba | H.K)$ will be very small, and the second addend on the right hand side of (2) can be

neglected (since $1 - x$ is close to zero). Now we come to the left hand side. Regardless of whether we observe black ravens or white shoes, a single observation of either object will not impose major changes on the posterior probability of H . This transfers to the *posterior odds* of H after observing $Ba.Ra$ or $\neg Ba.\neg Ra$, respectively. Thus, the quotient of those posterior odds will be close to 1 — even more close than $x = P(Ba \mid Ra.\neg H.K)$. And by this line of reasoning, we have established (2) and the last of Fitelson and Hawthorne’s premises. Thus, their argument is not only valid, but also conclusive.

Of course, it is still possible to doubt one of the plausibility arguments in the previous paragraphs. But I think they are cogent enough to put the burden of proof to those who doubt Fitelson and Hawthorne’s comparative solution. Moreover, the elegance of their proof deserves high praise, and since they use clear-cut assumptions, their analysis directly points out the points of disagreement between defenders and critics of their solution. Furthermore, they do not rely on the independence claim (IA) or variants thereof.

6 SUMMARY

The first part of this article has described and reviewed Hempel’s theory of confirmation and his analysis of the paradoxes of confirmation. Hempel’s approach to modeling confirmation departs from Carnap’s probabilistic approach: he decides to lay the qualitative foundations first by formulating general adequacy constraints that any account of confirmation has to satisfy. Hempel’s qualitative account of confirmation breaks with the classical hypothetico-deductive approach and proposes the *satisfaction criterion*: the restriction of a hypothesis to a specified object domain has to be entailed by the evidence. The criterion, however, has several shortcomings, some of them of technical nature, others being connected to the failure to account for confirmation by successful prediction. One of the most severe objections contends, however, that the satisfaction criterion is often monotone with respect to the background knowledge and thus unable to deal with the *paradoxes of confirmation*. On the one hand, Hempel has convincingly argued that the paradoxes rest on a psychological illusion, due to the tacit introduction of additional background knowledge. But on the other hand, his own criterion of confirmation neglects that insight and therefore fails to remove the paradox.

The second part of the article focuses on recent attempts to solve the paradoxes in the framework of Bayesian confirmation theory. While Hempel was probably right that the qualitative version of the paradox was just a *scheinproblem*, there are comparative and quantitative versions of the paradox, too. To vindicate these intuitions in a probabilistic framework has proved to be a tough task. By their [2009] result, Fitelson and Hawthorne solve the comparative problem and give some reasons for optimism. But so far, the quantitative problem remains unsolved. Even more embarrassing, I have argued that there are principal problems that impair a sufficiently general resolution of the paradoxes of confirmation. The conclusion which I draw — scepticism towards quantitative solutions of the paradox — is

somewhat atypical because most contributions to the literature either propose a solution or suggest to replace a previous attempt by a novel and better one.³⁴ But the longstanding history of the paradox indicates that it will be hard to overcome.

ACKNOWLEDGEMENTS

I would like to thank Andreas Bartels, Branden Fitelson, Stephan Hartmann, James Hawthorne, Franz Huber, Kevin Korb, and Jacob Rosenthal for their incredibly helpful advice and criticism.

BIBLIOGRAPHY

- [Agassi, 1959] J. Agassi. Corroboration versus Induction, *British Journal for the Philosophy of Science*, 9:311-317, 1959.
- [Alexander, 1958] H. G. Alexander. The Paradoxes of Confirmation, *British Journal for the Philosophy of Science*, 9:227-233, 1958
- [Alexander, 1959] H. G. Alexander. The Paradoxes of Confirmation — A Reply to Dr Agassi, *British Journal for the Philosophy of Science*, 10:229-234, 1959.
- [Black, 1966] M. Black. Notes on the 'paradoxes of confirmation'. In Jaakko Hintikka and Patrick Suppes, eds., *Aspects of Inductive Logic*, pp. 175-197, North-Holland, Amsterdam, 1966.
- [Carnap, 1950] R. Carnap. *Logical Foundations of Probability*, The University of Chicago Press, Chicago, 1950.
- [Carnap, 1952] R. Carnap. *The Continuum of Inductive Methods*, The University of Chicago Press, Chicago, 1952.
- [Dietrich and Moretti, 2005] F. Dietrich and L. Moretti. On Coherent Sets and the Transmission of Confirmation, *Philosophy of Science*, 72:403-424, 2005.
- [Duhem, 1914] P. Duhem. *La Théorie Physique: Son Objet, Sa Structure*. 1914. Second edition, reprinted in 1981 by J. Vrin, Paris.
- [Earman, 1992] J. Earman. *Bayes or Bust?* The MIT Press, Cambridge/MA, 1992.
- [Earman and Salmon, 1992] J. Earman and W. Salmon. The Confirmation of Scientific Hypotheses. In Merrilee H. Salmon, ed., *Introduction to the Philosophy of Science*, pp. 42-103. Hackett, Indianapolis, 1992.
- [Fitelson, 2001] B. Fitelson. A Bayesian Account of Independent Evidence with Applications, *Philosophy of Science*, 68:S123-S140, 2001.
- [Fitelson, 2006] B. Fitelson. The Paradox of Confirmation, *Philosophy Compass*, 1:95-113, 2006
- [Fitelson, 2009] B. Fitelson. The Wason Task(s) and the Paradox of Confirmation, *Synthese*, 2009.
- [Fitelson and Hawthorne, 2009] B. Fitelson and J. Hawthorne. How Bayesian Confirmation Theory Handles the Paradox of the Ravens. In Ellery Eells and James Fetzer, eds., *Probability in Science*, Open Court, Chicago, 2009.
- [Friedman, 1999] M. Friedman. *Reconsidering Logical Positivism*, Cambridge University Press, Cambridge, 1999.
- [Gaifman, 1979] H. Gaifman. Subjective Probability, Natural Predicates and Hempel's Ravens, *Erkenntnis*, 21:105-147, 1979.
- [Gemes, 1993] K. Gemes. Hypothetico-Deductivism, Content and the Natural Axiomatization of Theories, *Philosophy of Science*, 60:477-487, 1993.
- [Gemes, 2006] K. Gemes. Content and Watkins' Account of Natural Axiomatizations, *dialectica*, 60:85-92, 2006.
- [Glymour, 1980] C. Glymour. *Theory and Evidence*, Princeton University Press, Princeton, 1980.
- [Good, 1960] I. J. Good. The Paradox of Confirmation, *British Journal for the Philosophy of Science*, 11:145-149, 1960.

³⁴[Korb, 1994] and [Vranas, 2004] are notable exceptions.

- [Good, 1961] I. J. Good. The Paradox of Confirmation (II), *British Journal for the Philosophy of Science*, 12:63-64, 1961.
- [Good, 1967] I. J. Good. The White Shoe is a Red Herring, *British Journal for the Philosophy of Science*, 17:322, 1967.
- [Good, 1968] I. J. Good. The White Shoe qua Herring is Pink, *British Journal for the Philosophy of Science*, 19:156-157, 1968.
- [Goodman, 1983] N. Goodman. *Fact, Fiction and Forecast*, Fourth Edition. Harvard University Press, Oxford, 1983.
- [Hempel, 1943] C. G. Hempel. A Purely Syntactical Definition of Confirmation, *Journal of Symbolic Logic*, 8: 122-143, 1943.
- [Hempel, 1965] C. G. Hempel. Studies in the Logic of Confirmation, *Aspects of Scientific Explanation*, 3-51, 1965. The Free Press, New York. Reprint from *Mind* 54, 1945
- [Hempel, 1967] C. G. Hempel. The White Shoe: No Red Herring, *British Journal for the Philosophy of Science*, 18:239-240, 1967.
- [Hosiasson, 1940] J. Hosiasson-Lindenbaum. On Confirmation, *Journal of Symbolic Logic*, 5:133-148, 1940.
- [Horwich, 1982] P. Horwich. *Probability and Evidence*, Cambridge University Press, Cambridge, 1982.
- [Howson and Urbach, 1993] C. Howson and P. Urbach. *Scientific Reasoning: The Bayesian Approach*. Second Edition. Open Court, La Salle, 1993.
- [Huber, 2008] F. Huber. Hempel's Logic of Confirmation, *Philosophical Studies*, 139:181-189, 2008.
- [Humburg, 1986] J. Humburg. The Solution of Hempel's Raven Paradox in Rudolf Carnap's System of Inductive Logic, *Erkenntnis*, 24:57-72, 1986.
- [Korb, 1994] K. B. Korb. Infinitely many solutions of Hempel's paradox. *Theoretical Aspects Of Rationality and Knowledge — Proceedings of the 5th conference on Theoretical aspects of reasoning about knowledge*. pp. 138-149. San Francisco: Morgan Kaufmann Publishers, 1994.
- [Mackie, 1963] J. L. Mackie. The Paradox of Confirmation, *British Journal for the Philosophy of Science*, 13:265-276, 1963.
- [Maher, 1999] P. Maher. Inductive Logic and the Ravens Paradox, *Philosophy of Science*, 66:50-70, 1999.
- [Maher, (2004)] P. Maher. Probability Captures the Logic of Confirmation. In Christopher Hitchcock, ed., *Contemporary Debates in the Philosophy of Science*, pp. 69-93. Blackwell, Oxford, 2004.
- [Musgrave, 2009] A. Musgrave. Popper and Hypothetico-Deductivism. In this volume, 2009.
- [Nicod, 1961] J. Nicod. *Le Problème Logique de l'Induction*. Paris: Presses Universitaires de France. Originally published in 1925 (Paris: Alcan), 1961.
- [Popper, 1963] K. R. Popper. *Conjectures and Refutations: The Growth of Scientific Knowledge*, Routledge, London, 1963.
- [Royall, 1997] R. Royall. *Statistical Evidence: A Likelihood Paradigm*, Chapman & Hall, London, 1997.
- [Schurz, 1991] G. Schurz. Relevant Deduction, *Erkenntnis*, 35:391-437, 1991.
- [Suppes, 1969] P. Suppes. Models of Data. I P. Suppes, ed., *Studies in the Methodology and Foundations of Science. Selected Papers from 1951 to 1969*, pp. 24-35. Reidel, Dordrecht, 1969. Originally published in Ernest Nagel, Patrick Suppes and Alfred Tarski (eds.): "Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress". Stanford: Stanford University Press, 252-261, 1962.
- [Swinburne, 1971] R. Swinburne. The Paradoxes of Confirmation — A Survey, *American Philosophical Quarterly*, 8:318-330, 1971.
- [Uebel, 2006] T. Uebel. The Vienna Circle, 2006.
- [Vincent, 1964] D. H. Vincent. The Paradoxes of Confirmation, *Mind*, 73:273-279, 1964.
- [vonWright, 1966] G. H. von Wright. The Paradoxes of Confirmation. In Jaakko Hintikka and Patrick Suppes, eds., *Aspects of Inductive Logic*, pp. 208-218. North-Holland, Amsterdam, 1966.
- [Vranas, 2004] P. Vranas. Hempel's Raven Paradox: A Lacuna in the Standard Bayesian Solution, *British Journal for the Philosophy of Science*, 55:545-560, 2004.
- [Wasonand Shapiro, 1971] P. C. Wason and D. Shapiro. Natural and contrived evidence in a reasoning problem, *Quarterly Journal of Experimental Psychology*, 23:63-71, 1971.

- [Watkins, 1957] J. W. N. Watkins. Between Analytical and Empirical, *Philosophy*, 33:112-131, 1957.
- [Watkins, 1960] J. W. N. Watkins. Confirmation without Background Knowledge, *British Journal for the Philosophy of Science*, 10:318-320, 1960.
- [Weisberg, 2009] J. Weisberg. Varieties of Bayesianism, this volume, 2009.
- [Woodward, 1985] J. Woodward. Critical Review: Horwich on the Ravens, Projectability and Induction, *Philosophical Studies*, 47:409-428, 1985
- [Zabell, 2009] S. Zabell. Carnap and the Logic of Induction, this volume, 2009.