**Concerning Reichenbach's Neglected Remarks on the Ravens Paradox**

John T. Roberts

Full version of May 24, 2011

Abstract: Reichenbach's brief remarks on the Ravens Paradox reward attention in (at least) two ways. First, Reichenbach offers a resolution of the Ravens Paradox that seems to have gone unnoticed in most of the contemporary literature. This resolution is quite different from the familiar ones. I argue that it is interesting and worth taking seriously. Reichenbach presented this resolution in the context of his own theory of confirmation, but I show that the key move is independent of that context and can be integrated with a standard Bayesian account of confirmation or an explanationist approach to confirmation. Second, one of Reichenbach's remarks about his resolution points the way to an interesting rebuttal to an otherwise plausible objection to the standard Bayesian resolution of the paradox.

In his 1949 *Theory of Probability* and his 1954 *Nomological Statements and Admissible Operations* (reprinted in 1976 as *Laws, Modalities and Counterfactuals*) Reichenbach made some brief and somewhat cryptic remarks on Hempel's (1945) paradox of confirmation, also known as the Ravens Paradox. These remarks point the way to an interesting resolution of the Ravens Paradox that is very different from any of the familiar resolutions found in more recent literature. Reichenbach himself presented the resolution in the context of his own distinctive approach to the theory of evidence, but I will argue that the key idea can be used by advocates of other approaches as well. It can be adopted within a Bayesian framework, leading to a Bayesian resolution of the paradox that contrasts sharply with the standard Bayesian resolution; it can also be adopted within a non-Bayesian, explanationist approach to evidence and confirmation, and when it is it provides a neat solution to a problem otherwise faced by such an approach. Moreover, Reichenbach's remarks also point to an interesting fact about the Ravens Paradox: It shows that a certain advantage which some alternatives to the standard Bayesian resolution of the paradox appear to have over that resolution is illusory.

Before discussing Reichenbach's remarks, I will do some stage-setting. I will briefly review the Ravens Paradox and its standard Bayesian resolution, as well as an advantage sometimes claimed for other resolutions of the paradox.

**1. The Paradox Restated**

We assume that unless something strange is going on, the discovery of a positive instance of a universal generalization confirms that generalization.[1] There are some unusual circumstances in which the discovery of a positive instance fails to confirm the generalization, or even disconfirms it.[2] But those are the odd cases; in the typical case, in the absence of particular kinds of defeating conditions, we think that discoveries of positive instances of a generalization confirm that generalization.[3] Call this the *Instance-Confirmation Assumption*. We also assume that whatever confirms one hypothesis also confirms any hypothesis logically equivalent to it; call this the *Equivalence Condition*. The Instance-Confirmation Assumption implies that (so long as nothing funny is going on) the discovery of a non-black non-raven – such as a green beer bottle – confirms the

---

[1] Strictly speaking, we should say it is a *statement* or *proposition* that confirms the generalization – the statement that a certain object is a black raven, for instance, or the statement that the first raven we sampled was black, or that the $n^{th}$ bird we sampled was a black raven. I won't bother to be strict about this point in the text when not doing so should create no misunderstanding.

[2] Some imaginative examples are found in Good (1967) and (1968).

[3] Most presentations of the paradox take one premise to be *Nicod's criterion,* which says that a positive instance of a universal generalization always confirms it. I prefer instead to use the hedged premise in the text, because using the full-strength Nicod's criterion makes it seem too easy to escape from the paradox: We know there are lots of counterexamples to Nicod's criterion, so problem solved, right? (The discovery of a carpenter ant in Orange County, one centimeter from the county line, disconfirms the hypothesis that all carpenter ants are inside Orange County; in a context where the background knowledge includes the proposition that there are probably no ravens at all but if there are any at all then there are most likely ravens of many different colors, the discovery of a single black raven disconfirms the hypothesis that all ravens are black.) The trouble is that all those familiar counterexamples are explicable in terms of identifiable things going on in the situation that seem to have no analogue in the original ravens case. Setting up the paradox in the way I have just done does not result in a crisp, logically valid deduction, which is a shame; however, my version avoids misleadingly suggesting that finding any reason to doubt one of the premises would suffice to solve the problem.

hypothesis that all non-black things are non-ravens. By the Equivalence Condition, it follows that the same piece of evidence confirms the hypothesis that all ravens are black. The conclusion seems delightfully paradoxical, since it seems to permit us to conduct ornithological research at the pub.[4]

## 2. The Standard Bayesian Resolution of the Paradox

The standard Bayesian resolution[5] comes in many versions, but they all have the following things in common: They accept that the paradox offends some of our pre-theoretical judgments about what does and does not confirm a hypothesis that takes the logical form of a strict universal generalization, $\forall x(Fx \supset Gx)$. They embrace the conclusion that Hempel's paradoxical conclusion is in fact true—that is, the discovery of a non-G non-F really does confirm the hypothesis that $\forall x(Fx \supset Gx)$. Finally, they endeavor to blunt the sting of this paradoxical conclusion in two ways: first by arguing that under certain conditions, the degree of confirmation that the discovery of a non-G non-F confers on the hypothesis is miniscule, so our pre-theoretical judgments are not far off when they tell us that such a discovery is confirmationally irrelevant; second, they provide an evidently compelling argument that, in spite of first-blush appearances, there are excellent reasons to think that a non-G non-F really does confirm the hypothesis that all Fs are Gs (if only by a tiny amount), so that we should be willing to have our pre-theoretical judgments corrected on this point.

For Bayesians, evidence E confirms hypothesis T just in case conditionalizing on E raises our credence for T: $Cr(T|E) > Cr(T)$. Degree of confirmation is the degree to which E raises the credence of T; there are several ways of measuring it, and Bayesians

[4] Hempel (1945), pp. 14-15. Hempel presents another, related paradox as well: The Instance-Confirmation Condition and the Equivalence Condition jointly entail that the discovery of a black non-raven confirms the hypothesis that all ravens are black. In this paper I discuss only the first paradox.

[5] First published in Hosiasson-Lindenbaum (1940); influentially presented in Mackie (1963) and Howson and Urbach (1993), 126-129; for an overview of the history, see Fitelson (2006) and for a sophisticated contemporary version of the solution with some important technical improvements over previous versions see Fitelson and Hawthorne (2006).

are not all in agreement about which way is most appropriate. Two common measures are:

*The ratio measure:* The degree to which E confirms T is proportional to:

$$Cr(T|E) / Cr(T)$$

*The likelihood ratio measure:* The degree to which E confirms T is proportional to:

$$Cr(E|T) / Cr(E|\sim T)$$

Other measures are possible as well; for example, the logarithm of the ratio measure, or the logarithm of the likelihood ratio.

One crucial plank of the standard Bayesian resolution of the Ravens Paradox is the claim that, under a certain condition, the degree to which the discovery of a non-black non-raven confirms the hypothesis that all ravens are black is only miniscule. This is supposed to explain away, or perhaps vindicate, or perhaps simply exculpate our common intuitive judgment that such a discovery does not confirm the hypothesis at all; a miniscule degree of confirmation is easily confused with a complete lack of confirmation. This plank is crucial because without it, our sense of the paradoxical here is not diffused; we might understand why there are good theoretical reasons to believe Hempel's paradoxical conclusion, but we will be unable to make sense of why that conclusion seems so obviously false.

The "certain condition" under which the degree of confirmation is miniscule is that before we examine our evidence, we find it more likely that a randomly sampled object will be non-black than that it will be a raven. (Or as some authors put it: We know that there are more non-black objects in the universe than there are ravens.)[6] Let's

[6] As Vranas (2004) points out, most published versions of the standard solution also depend on further assumptions to the effect that whether a given individual is a raven, or is black, is statistically independent of whether all ravens are black. Vranas also points out that these assumptions are very hard to justify. Fitelson and Hawthorne (2006) argue that the standard Bayesian solution can dispense with

see why this condition must be met, given either of the measures of confirmation just mentioned.

Suppose first that we adopt the ratio measure. Suppose that we are sampling from the population of non-black things, and let ~R be the proposition that what we find is a non-raven, and let H be the hypothesis that every raven is black. On the ratio measure, a value of 1 means that evidence E is confirmationally neutral with respect to hypothesis T; a value greater than 1 means that E is positively confirmationally relevant to T, and the greater the value, the more support it gives T. Given this measure, what it takes for ~R to confirm H, though only to a miniscule degree, is for this ratio:

$$
\begin{aligned}
&Cr(H|{\sim}R) / Cr(H) \\
&= Cr({\sim}R|H) / Cr({\sim}R) \\
&= 1/Cr({\sim}R)
\end{aligned}
$$

to have a value that is greater than 1, but only slightly. Therefore, Cr(~R) must be less than 1, but very close to 1. That is, when sampling from a population of non-black things, we would have to be very confident that we are not going to find a raven. And this must be so before we have begun finding evidence for the hypothesis that all ravens are black.[7] How can we be so sure that when we sample the non-black things, we are extremely unlikely to find a raven, even when, for all we know, it could well be that all or nearly all of the ravens are non-black? We must start out very confident that the other non-black things greatly outnumber the non-black ravens, no matter what proportion of the ravens are non-black. This can be so only if we start out very confident that there are many more non-black things than there are ravens of any color.

---

these independence assumptions. They retain the assumption that the probability of being non-black is greater than that of being a raven, however.

[7] Note that even before we have observed a single Martian, or have any evidence whatsoever about their color, it seems to us that if we want to confirm "all Martian are green," the thing to do is to observe Martians and check for greenness, not to observe non-green things and check to see if they are non-Martians. Indeed, nothing essential to the paradox changes when we switch to an example such as this, in which we aren't even sure if the hypothesis in question is true. So a good solution to the paradox had better work even when we don't yet have much confidence at all in the truth of the hypothesis in question.

On the log-ratio measure, positive but miniscule confirmation means that the logarithm of the ratio we were just considering must be positive but tiny. This implies that the logarithm of Cr(~R) must be negative but very close to zero; equivalently, Cr(~R) must be less than 1, but very close to 1. So the same argument applies, and we reach the same conclusion.

Now consider the likelihood ratio:

$$Cr(E|T)/Cr(E|{\sim}T)$$

Here again, a value of 1 indicates evidential neutrality, while a ratio greater than 1 means that E is positively evidentially relevant to T, and the greater the value the greater the support it gives T. So in order for ~R to confirm H, but only to a miniscule degree, the ratio:

$$\frac{Cr(\sim R \mid H)}{Cr(\sim R \mid \sim H)} = \frac{1}{Cr(\sim R \mid \sim H)}$$

must be greater than 1, but only slightly; so, Cr(~R|~H) must be less than 1, but only slightly. Again, this must be so before we gain any significant information about what proportion of the ravens are black.[8] Conditional on the hypothesis that at least some ravens are not black, and without any further information about the proportion of ravens that are, we must be very confident that when we randomly sample from the non-black things, we will find a non-raven. For the same reason as before, it follows that the standard Bayesian resolution works only if we start out quite confident that there are far more non-black things than there are ravens.

Finally, should we adopt the log-likelihood-ratio measure, then we find that in order for the degree of confirmation to be positive but minimal, the logarithm of the ratio 1/Cr(~R|~H) must be greater than zero but very small; this implies that Cr(~R|~H) must be less than 1 but only slightly so; the rest of the argument continues as in the previous paragraph.

[8] See the preceding note.

So on any of the four measures of confirmation we have considered, the degree of confirmation conferred on the hypothesis that all ravens are black by the discovery of a non-black non-raven is miniscule only on the condition that we start out very confident that there are more non-black things than there are ravens. Therefore, the standard Bayesian resolution of the paradox is satisfying only if this condition is met.

### 3. A Disadvantage of the Standard Bayesian Resolution?

There exist a number of non-Bayesian resolutions of the Ravens Paradox which do not seem to depend on that condition. These include the resolutions due to Scheffler (1981, pp. 286-91) following Goodman (1945, pp. 71-2), Armstrong (1983, pp. 41-6, pp. 102-3), Lange (2000, pp. 139-142), Lipton (2004, Chapter 6), and Lipton (2007). Moreover, it appears at first that Reichenbach's resolution is independent of this assumption. (Things are more complicated than they appear, though; see below, section 9.) Is this an advantage that those other resolutions enjoy over the standard Bayesian one? It is just in case the paradoxical conclusion still seems paradoxical even in cases where that assumption is false (as Hempel (1945, pp. 20-21, note 25) pointed out). Are there such cases?

Scheffler (1981, p. 283) argues that there are. One example he offers is the hypothesis "All molecules are inanimate." Suppose that we attempted to confirm this hypothesis by noting that it is logically equivalent to "All animate things are non-molecules" (understanding "non-molecule" to cover anything that is not identical to some particular molecule, including things that are composed of many molecules), and then inspecting many and varied animate things (dogs, cats, begonias, etc.), noting that each is a non-molecule. Surely our intuitions balk at this as firmly as they balk at the notion of confirming the hypothesis that all ravens are black by discovering non-black non-ravens. Yet, we know that there are far more molecules than there are animate things, so this is a case where the assumption that the standard Bayesian resolution depends on is false. Scheffler in effect argues that the paradox is just as palpable in the case of the molecules as it is in the case of the ravens, and indeed it appears to work in exactly the same way, but the standard Bayesian resolution is worthless in the case of the molecules, so we should not consider it to be the complete solution to the problem.

The stage-setting is complete; now let's turn to what Reichenbach said about the paradox.

## 4. Reichenbach's Remarks on the Paradox

In his *Theory of Probability*, Reichenbach presents the paradox as hinging on the logical equivalence of propositions he calls (1) and (2):

(1) $(\forall x)(Fx \supset Gx)$

(2) $(\forall x)(\sim Gx \supset \sim Fx)$[9]

Reichenbach writes:

> The paradox seems to be unsolvable within the theory of confirmation [i.e., qualitative confirmation theory]. It disappears, however, as soon as probabilities are introduced. For a probability implication, even of degree 1, contraposition does not hold, and thus the two forms corresponding to (1) and (2)
>
> (3) $P(Gx|Fx) = 1$
>
> (4) $P(\sim Fx|\sim Gx) = 1$
>
> are not equivalent. This fact can also be made clear as follows. If (3) were manifestly false and we had only a probability implication of a low degree [i.e. if $P(Gx|Fx)$ were low], (4) might remain virtually true. … Consequently, if the truth of (4) is established to a high degree of [epistemic] probability, (3) need not be true, and we cannot use an establishment of (4) as a proof for the validity of (3).

In his *Nomological Statements and Admissible Operations*,[10] Reichenbach makes the stronger claim that the same point holds even if the two conditional probabilities

[9] Reichenbach 1949, 434-5. Here and throughout this section, I have translated from Reichenbach's notation into more standard contemporary notation.
[10] Reprinted as Reichenbach 1976.

occurring in (3) and (4) are only approximately equal to 1. There he discusses equations he calls (158) and (159):

(158) $P(Bx|Ax) \approx 1$[11]

(159) $P({\sim}Ax|{\sim}Bx) \approx 1$

These propositions are not equivalent, as an example due to Wesley Salmon shows: Let Ax stand for ‘x is a philosopher,’ let Bx stand for ‘x is unkind,’ and interpret the probability function P as standing for relative frequencies; then (158) says that approximately 100% of philosophers are unkind, which is plainly false, while (159) says that approximately 100% of kind people are non-philosophers, which is surely true, just because approximately 100% of *people* are non-philosophers. Reichenbach writes:

> Relation (159) cannot be derived from (158) unless further conditions are specified. When we put
>
> (160) $P(Bx|Ax) = 1 - d,\ P({\sim}Ax|{\sim}Bx) = 1 - d'$
>
> …we find[12]:
>
> (161) $P({\sim}Bx|Ax) / P(Ax|{\sim}Bx) = d/d' = (1 - P(Bx))/P(Ax)$
>
> Therefore $d' \leq d$ if and only if $P(Ax) \leq 1 - P(Bx)$, or
>
> (162) $P(Ax) + P(Bx) \leq 1$

[11] We can understand $P(Bx|Ax)$ to mean the probability that a randomly selected object with property A will have the property B; Reichenbach himself, of course, understands these probabilities as relative frequencies.

[12] Reichenbach’s argument depends on the following theorem of probability:

$$P(X|Y)/P(Y|X) = P(X)/P(Y)$$

Plugging in ~Bx for X and Ax for Y gives us:

$$d/d' = P({\sim}Bx|Ax)/P(Ax|{\sim}Bx) = P({\sim}Bx)/P(Ax) = (1 - P(Bx))/P(Ax).$$

> Relation (162) formulates the condition on which we can go from (158) to (159) while remaining within the same small deviation d from the value 1. (Reichenbach 1976, 130)

(162) implies that (158) is a sufficient condition for (159). However, as Reichenbach goes on to note, (162) does not imply that (159) is a sufficient condition for (158). To infer from (159) to (158), what we need to know is that $d \leq d'$, that is:

(163) $P(Ax) + P(Bx) \geq 1$

> Conditions (162) and (163) are compatible only for the special case of the equality sign [that is, $P(Ax) + P(Bx) = 1$]. In general, therefore, we can proceed only in one direction. For instance, if we know that (162) holds, a proof of (158) is a proof of (159), but not vice versa. In the usual application of inductive verification, the terms 'A' and 'B' are so defined that (162) is satisfied. (ibid, 131)

For instance, in one typical application 'Ax' stands for 'x is a raven' and 'Bx' for 'x is black'; pretty obviously[13], $P(Ax) < ½$ and $P(Bx) < ½$, so condition (162) is satisfied, but not (163). Reichenbach goes on:

> This explains why confirming evidence for (158) is also confirming evidence for (159), whereas confirming evidence for (159) is not confirming evidence for (158). These considerations supply the answer to a so-called paradox of confirmation pointed out by C. Hempel. (ibid)

But *how* do these considerations supply the answer to the paradox?

---

[13] Assuming, with Reichenbach, that these probabilities are to be understood as frequencies, and assuming that both ravens and black things account for minorities of the objects in the domain.

Reichenbach has shifted attention from hypotheses of the form $\forall x(Ax \supset Bx)$ to hypotheses of the form $P(Bx|Ax) = 1$ or $P(Bx|Ax) \approx 1$; as we have seen, these conditional probability statements do not contrapose in the way that universally quantified material conditionals do. But how does this help with the paradox of the ravens? Why isn't this just a change of subject?

## 5. Reichenbach's Key Move

Here is what I think Reichenbach has in mind: When we randomly sample several ravens, find them all to be black, and inductively reach a conclusion about the entire population of ravens, we are inferring from a sample statistic to a population statistic; in particular, we are inferring from the frequency of blackness in our sample to the frequency of blackness in the population from which our sample was drawn. Of course, it is a mistake to infer that the frequency in the population exactly matches the frequency in our sample; the most we are ever entitled to infer is that the sample frequency is a good approximation to the population frequency. So, when we have examined several ravens and found them all to be black, what we may conclude with high probability is that approximately all ravens are black; each additional black raven discovered, while no non-black ones are found, makes that probability a bit higher, but does not change the hypothesis thus supported from an 'approximately-all' hypothesis to a 'precisely-all' hypothesis. So it is perhaps a mistake to say that the discovery of black ravens serves as evidence for the hypothesis that, literally speaking, every single raven in the world is black; it would be better to say that it is evidence for the hypothesis that *about* all ravens are black.[14] [15] I say "perhaps" here because it is plausible that in everyday language, statements of the form "All ravens are black" really express the weaker, approximative relation $P(\text{Black}(x)|\text{Raven}(x)) \approx 1$, rather than the universally quantified $\forall x(\text{Raven}(x) \supset \text{Black}(x))$. (After all, students regularly accept "All ravens are

---

[14] In the case where there are finitely many ravens, of course, this means that the hypothesis that is supported is $P(\text{Black}(x)|\text{Raven}(x)) \approx 1$, not $P(\text{Black}(x)|\text{Raven}(x)) = 1$. In the case where there are infinitely many ravens, however, the hypothesis that $P(\text{Black}(x)|\text{Raven}(x)) = 1$ is a version of the approximative hypothesis, rather than the hypothesis that each and every raven is black.

[15] Presumably, it also confirms the relation $P(\text{Raven}(x)|\text{Black}(x)) \approx 1$.

black" as a perfectly fine example with no balking at all, even though they must know that there are such things as albino ravens.) So, the intuition that leads us to affirm the Instance Confirmation Assumption is really tracking, not the fact (or alleged fact) that is strictly and literally expressed by the formulation of it that I gave above, but instead a fact we might call the Approximative Instance-Confirmation Assumption: *In typical cases, in the absence of particular kinds of defeating conditions, discoveries of positive instances of a universal generalization confirm the approximate truth of that generalization.*

If this is right, then as Reichenbach says, it resolves the paradox: The discovery of a non-black non-raven confirms not the hypothesis that all non-ravens are non-black, but instead the hypothesis that *about* all non-ravens are non-black, i.e. all-or-nearly-all non-ravens are non-black. Now the Equivalence Condition gains no purchase: The proposition that *about all non-ravens are non-black* is logically independent of the hypothesis that *about all ravens are black*, so we do not have here any confirmation for the latter hypothesis. Green beer bottles do not, after all, serve as valuable ornithological evidence.

Note that this way of resolving the paradox works very differently from the standard Bayesian resolution: It takes the paradoxical conclusion to be true, and 'resolves' the paradox by showing that there is an important sense in which our intuitions about the case are approximately true (viz., that a non-black non-raven confirms the hypothesis that all ravens are black only to a very tiny degree) even though they are not exactly true. By contrast, Reichenbach's resolution assumes that our intuitions about the case are correct—the discovery of a non-black non-raven is no evidence at all for the hypothesis we typically express with the words "All ravens are black"—and it explains away the appearance of paradox by pointing out that this appearance was generated by an inappropriate formalization of that hypothesis: It is more properly rendered as $P(Black(x)|Raven(x)) \approx 1$ than as $\forall x(Raven(x) \supset Black(x))$.

Reichenbach incorporates this insight about the ravens paradox into his own theory of scientific evidence, but it is clear that the basic idea here is independent of that theory. Below, I will show how to integrate it into a Bayesian approach to confirmation, and then show how to integrate it into a non-Bayesian abductive framework. First, though, I will consider and respond to an important objection.

## 5. An Objection and a Reply

The Reichenbachian resolution doesn't work if the discovery of a black raven confirms the hypothesis that (exactly) all ravens are black. (For if it does confirm that hypothesis, then presumably the discovery a non-black non-raven confirms the hypothesis that all non-ravens are non-black, so unless we want to give up the Equivalence Condition we are stuck with the paradoxical conclusion.) But it seems obvious to many of us that such a discovery can confirm that hypothesis. If we understand the verb 'confirm' to mean 'increase our rational confidence in,' then even if the discovery does confirm the hypothesis that about all ravens are black, it can simultaneously confirm the hypothesis that exactly all ravens are black. What's more, it seems obvious that we have confirmed many hypotheses that do take the form of universally quantified conditionals—for example, *all material particles travel more slowly than light*, *energy is always conserved in closed systems*, etc. So isn't it unreasonable to suppose that discovering instances of a universal generalization cannot confirm the (not-merely-approximate) truth of that very generalization?

Furthermore, it might seem that on a Bayesian account of confirmation, the discovery of a positive instance of a universal generalization simply must confirm it. Assume that E is a positive instance of G; then ~E entails ~G, so the discovery that E is true is the discovery that one of the ways in which G might have been false fails to transpire; this must raise our credence in G, mustn't it? Well, not if our prior credence for G was zero; in that case, conditionalizing on E (or anything else) must result in a posterior of zero as well. When we consider that a universal generalization $\forall x(Ax \supset Bx)$ entails a hypothesis to the effect that a certain real-valued parameter—namely, the frequency of B among the As—has a certain definite value, we can see a rationale for thinking that it might not be at all unreasonable for G's prior to be zero. This shows that Bayesianism by itself does not imply that the discovery of a positive instance of a universal generalization must confirm it, in the sense of raising its probability.

Nevertheless, it might seem obvious on other grounds that in many cases, the discovery of a positive instance does confirm a generalization. Let's think more carefully about a particular example: Perhaps we now have excellent reason to believe *that no massive body ever travels faster than the speed of light*, which can be written in the form of a strict universal generalization. What is our evidence for this claim? It certainly does not consist solely of lots and lots of particular observations of massive bodies failing to travel faster than the speed of light. Our evidence for this claim does not support it via enumerative induction or instantial confirmation; it supports it via a sophisticated theoretical inference, involving the Lorentz transformations and special-relativistic kinematics and dynamics. Those hypotheses, in turn, are supported by theoretical inferences involving, among other things, Maxwell's equations and symmetry considerations, all of which are in turn supported by further theoretical considerations. To be sure, empirical evidence plays an important role in this process, and at each stage the claims supported by theoretical considerations have also been put to direct tests in which positive instances of them have been confirmed. But the role played by such tests is clearly not that of providing instances to be used in enumerative induction. (This is made clear by, among other things, the fact that in tests of high-level theories, often only one or a few observations are enough to make scientists extremely confident of the claims under test. More and more repetitions of the Michelson-Morley experiment do not greatly increase our confidence in the Lorentz transformations.) So in cases where we think we have excellent evidence for the truth of some universal generalization that enjoys impressive theoretical support and has been subjected to stringent tests, it seems likely that for the most part its support does not come from enumerative induction or instantial confirmation, but instead form some more sophisticated sort of inference. (Alternatively, theoretical considerations might give it a non-zero prior, which in turn gets boosted by discovering instances; perhaps this is what happens in the speed-of-light case.) Adopting the Reichenbachian approach to the ravens paradox will therefore most likely not undermine our ability to make sense of why we think we have great evidence for many such generalizations.

It is a serious problem, though, if there are clear cases of strict universal generalizations that we have good reason to think strictly true, where these reasons do not come from theoretical considerations, but only from instantial confirmation or enumerative induction. Are there any such clear cases? "All humans are mortal" springs to mind—but surely this generalization does enjoy broad theoretical support, coming from many quarters of the life sciences. Ditto for "All pine trees spring from pine cones" and the like. "All cats meow," "All politicians kiss babies," "All Belgian beers are heavy and syrupy" – these all have some plausibility, and no doubt the reader can supply more.[16] But it's hard to tell what our evidence really supports in these cases. What exactly is the difference between this:

> Having good reason to believe that all Fs are Gs—though you know of course that as future data come in you might have to give this belief up,

and this:

> Having good reason to believe that either all or nearly all Fs are Gs—though you know of course that as future data come in you might have to give this belief up?

The line between these cases seems very fine. Unless we have a case where there is no doubt about which side of that line we are on, we do not have a clear counterexample to the Reichenbachian proposal. In itself this is not an argument in favor of the proposal, but it does seem to justify us in not proclaiming the proposal dead on arrival; we can keep playing with it to see if it has any impressive virtues that might justify us in accepting it.

## 7. Reichenbach's Resolution in a Bayesian Setting

Let $H_r$ be the hypothesis that the frequency of Gs among the Fs is r, where r is any

---

[16] Matt Kotzen suggested to me the examples of Goldbach's conjecture and of Fermat's last theorem before it was proven: These are universal generalizations, and many mathematicians feel that the failure to discover any counterexamples after prolonged attempts made it reasonable to believe that they were probably true. Marc Lange suggested to me the example of Balmer's hypothesis that there are lines in the spectrum of hydrogen with wavelengths equal to $Bm^2/(m^2 - n^2)$ where n is any integer greater than 1 and m is any integer greater than n; when Balmer discovered lines corresponding to the first several values of m and n, it might have been reasonable to consider that to be confirming evidence for the generalization that there are lines corresponding to *all* the possible values of m and n, even though at the time there was no theoretical reason to believe this. (Bohr supplied such a theoretical reason later with his model of the atom.) These are all impressive examples, but I am inclined to reply to them in the same way I am about to reply to the aforementioned examples in the text.

real number between 0 and 1 inclusive. Consider a rational agent who starts out with some prior credence distribution over the hypotheses $H_r$. What should this credence distribution look like? Note that r is a real-valued parameter. In the absence of any theoretical considerations that constrain its value in such a way as to make some finite or countable set of its possible values particularly likely[17], considerations of symmetry suggest that our agent should not assign a positive probability to any particular value, but instead should assign a probability density $\delta(r)$ to the possible values. (I do not say that this should be a uniform density. Perhaps it should be, but we don't need to assume that.)[18]

---

[17] E.g., theoretical considerations that favor its being a law that all Fs are Gs, or that its being a law that the chance of an F being a G is in some countable range of possible values. I allow for a countable set of values here is because of considerations such as the following: It is conceivable (if just barely) that there are theoretical reasons that make it seem plausible that r is equal to $1/2^n$ for some positive integer n, with lower values of n seeming more plausible than higher values. In this case, it might be possible for a reasonable agent to set her initial credence that r is equal to some such value to ½, and her credence that it is not to ½. Then she might set her credence that r = ½ to ¼, that r= ¼ to 1/8, that r = 1/8 to 1/16, and so on. Her remaining credence could then be distributed continuously and uniformly over the remaining values; every value of r between 0 and 1 that does not take the form $1/2^n$ would have a credence density of ½. So there are coherent ways to assign non-zero values to each of a countably infinite set of values. But if theoretical considerations favored the values in a set that was larger than countable, then obviously there would be no coherent way to assign all of its members a credence greater than 0.

[18] What if it is known that the number of Fs is finite? Then it is known that r is a rational number, and the rationals are of course a countable subset of the reals. So in this case, a particular countable subset of the possible values are particularly likely; does this mean that in this case our subject need not have a continuous credence distribution over the real interval [0, 1]? Well, there are two alternatives: She could assign non-zero credence to some finite number of particular hypotheses $H_r$ but not to the remaining infinite set of such hypotheses – which would obviously be grossly undermotivated by the background knowledge that r must be *some* rational number or other – or she could do something like this: set $Cr(H_r)$ for each rational r equal to $2^{-f(r)}$, where f(r) is some injection of the rationals into the natural numbers. It is, I submit, difficult to conceive of a situation in which it would be remotely reasonable for a subject to have such a credence distribution. Be that as it may, the resolution of the ravens paradox that I propose here depends on the assumption that we are not in a situation in which it would be reasonable to have such credences.

So our agent's prior credence in $H_1$ (like her credence for any particular $H_r$) is 0. So subsequent conditionalization will not raise it above zero. However, for any e < 1, her prior credence in the hypothesis $H_{[e, 1]}$, i.e. the hypothesis that the frequency r is in the interval [e, 1], is:

$$Cr(H_{[e, 1]}) = \int_e^1 \delta(r)dr$$

which will in general be greater than 0. Subsequent discovery of instances of F that are also instances of G will typically raise the credence of this hypothesis. For example, let $E_n$ be the evidence statement that n Fs have been sampled, by simple random sampling, and all have been found to be G. Then to find out how much the discovery of $E_n$ will raise $Cr(H_{[e, 1]})$, we need to calculate the ratio:

$$\frac{Cr(H_{[e,1]} \mid E_n)}{Cr(H_{[e,1]})} = \frac{Cr(E_n \mid H_{[e,1]})}{Cr(E_n)} = \frac{Cr(E_n \mid H_{[e,1]})}{Cr(E_n \mid H_{[0,1]})}$$

In the last expression, both numerator and denominator have the form $Cr(E_n \mid H_{[a,b]})$; how do we evaluate an expression of this form? Well, suppose that $H_{[a,b]}$ is true: Then the frequency r of Gs among the Fs has some value between a and b. Suppose that it has the value x: Then the likelihood we are interested in is the probability of n distinct, independent events each of which has probability x, so it is $x^n$. Now just integrate $x^n$ over all of its possible values between a and b. Our integrand must also include the probability density associated with each of these values. For simplicity's sake, let's assume that this density is uniform. This means we'll just need to multiply our integrand by a constant, which will be 1 divided by the size of the range over which we are integrating, (b-a). Thus:

$$Cr(E_n \mid H_{[a,b]}) = \int_a^b \frac{x^n}{b-a}dx$$

So:

$$\frac{Cr(E_n \mid H_{[e,1]})}{Cr(E_n \mid H_{[0,1]})} = \frac{\int_e^1 \frac{x^n}{(1-e)} dx}{\int_0^1 x^n dx}$$

$$= \frac{\frac{1}{(n+1)(1-e)} x^{n+1} \Big|_e^1}{\frac{1}{n+1} x^{n+1} \Big|_0^1} = \frac{1-e^{n+1}}{1-e}$$

This ratio behaves as we should expect it to. For example, let e = 0.9; then here are the values of the ratio for n = 1, 10, 100, 200: 1.9, 6.8619, 9.9976, 9.9999. These numbers asymptotically approach 10, which is what we should expect, since our prior probability for $H_{[.9, 1]}$ is 0.1 (recall that for simplicity we assumed a prior distribution uniformly distributed over the possible values for r ranging from 0 to 1), so as our posterior for $H_{[.9, 1]}$ converges to 1, this ratio converges to 10. Obviously, there is nothing special about the value e = 0.9 here; had we chosen any other small interval [e, 1], we would have found that as the number of observed Fs increased, with all observed Fs found to be black, the posterior probability that frequency r lies in the interval [e, 1] converged to 1. This interval must have some non-zero width; the probability that the frequency is precisely 1 always remains zero.

In this way, discovery of positives instances of "All Fs are Gs" will confirm the hypothesis that it is either true or approximately true (to any specific degree of approximation you like). But this is obviously not logically equivalent to the hypothesis that "All non-Gs are non-Fs" is either true or approximately true. And the discovery of Fs that are G will not in general confirm the latter hypothesis, just as the discovery of non-Gs that are not Fs will not in general confirm the former one. To see why, let's try working through a calculation to mirror the one we did before. Let $D_n$ be the evidence statement that n non-Gs have been selected by simple random sampling, and that all of them have been found to be non-Fs. Now the ratio we need to calculate is:

$$\frac{Cr(H_{[e,1]} \mid D_n)}{Cr(H_{[e,1]})} = \frac{Cr(D_n \mid H_{[e,1]})}{Cr(D_n)} = \frac{Cr(D_n \mid H_{[e,1]})}{Cr(D_n \mid H_{[0,1]})}$$

But this time, the likelihoods cannot be calculated in the same straightforward way we used last time. Suppose again that $H_{[a,b]}$ is true; then the frequency r of Gs among the Fs has some value between a and b. Suppose that it has the value x. But the probability we are interested in is not x; it is the probability of a randomly selected non-G's being a non-F, which need not bear any particular relation to that of a randomly selected F's being a G. Indeed, the assumptions we have made about the case radically underdetermine the probabilities we need to know now in order to continue with out calculation. If special additional conditions are imposed on the case, then it might be possible to derive the conclusion that the discovery of non-G non-Fs will tend to confirm the hypothesis that about all Fs are G, but this paradoxical conclusion simply does not follow from the way the case is standardly set up. In this way the paradox is avoided.

This Bayesian resolution requires an assumption that the prominent Bayesian resolutions to the paradox do not require, namely that prior credence is continuously distributed over the continuum of possible values of r, with no particular value receiving a positive credence (in the absence of any theoretical considerations that favor some countable set of possible values over the others). Hardcore subjectivist Bayesians will not find this plausible, but if we are willing to accept some objective constraints on priors other than probabilistic coherence, this constraint is not unmotivated and doesn't seem unduly strong.[19]

## 8. Reichenbach's Resolution in an Explanationist Setting

An explanationist approach to confirmation understands the confirmation relation as roughly the converse of the explanatory relation: To a first approximation, E confirms H insofar as H, were it true, would explain E.[20] There is a familiar explanationist strategy for dealing with the paradox that is quite different from the standard Bayesian strategy

[19] Strictly speaking, my solution imposes one more "objective" constraint on priors – the one mentioned at the end of note 7. Again, this one does not seem terribly restrictive, by the lights of someone who is not a fundamentalist subjectivist Bayesian.

[20] Depending on how the explanatory relation is understood, and depending on how this first approximation is corrected, this thought needn't conflict with Bayesianism, though it might.

considered above.[21] In this section I will argue that this explanationist strategy can be strengthened by allying it with the basic idea presented in section 5.

The familiar explanationist treatment of the paradox runs as follows. If it is in fact true that all ravens are black, then this is presumably a lawlike generalization, and its truth can explain (or play a central role in explaining) why particular ravens – including the ones which we have observed so far – are black. By contrast, if it is true that all non-black things are non-ravens, then this fact needn't be able to play any particular role in explaining why a given non-black thing is a non-raven. (For example, the explanation of the sun's non-ravenhood has to do with its having formed via a process in which raven gametes played no role; the sun's non-blackness will have nothing to do with it.) So the fact that all ravens are black does explain why the observed ravens have been black, therefore the latter can confirm the former; by contrast, the fact that all non-black things are non-ravens does not explain why the observed non-black things have been non-ravens, so the latter cannot confirm the former. So there is an important asymmetry between the case of using black ravens to confirm "All ravens are black" and using non-black non-ravens to confirm "All non-black things are non-ravens." Thus, an explanationist account of inductive support allows us to escape the paradox.

But things are not quite that easy. As Roger White (2005) has pointed out, the regularity that all Fs are Gs need not explain *why any particular F is G* in order to explain *why all of the observed Fs are G*. Suppose that we have been searching diligently for an F that isn't G, and we haven't been able to find one, despite our best efforts; why have we failed? Well, if it turns out that there aren't any Fs that aren't G – in other words, if all Fs are G – that would be an excellent explanation of our failure, *even if this fact fails to explain why any particular F is a G*. For example, suppose that every member of a certain club has red hair, not because members of that club dye their hair red or anything like that, but simply because only red-haired people are allowed to join. In that case, the fact that all members of the club are red-haired explains why *all of the members of the club that we have managed to locate and identify are red-haired*, even though for each particular member of the club, the fact that he or she is a member of the club fails to

[21] Found e.g. in Armstrong (1983), pp. 41-46, though Armstrong's presentation differs substantially from the one to follow.

explain why he or she has red hair. The point here that is relevant for us is that when our evidence is that we have observed many ravens and all of them have been black, our hypothesis H can (if true) explain E without explaining why a single particular raven is black. Similarly, if our evidence is that we have searched diligently for non-black things in a variety of locations and conditions, and so far all of them have been found to be non-ravens, a hypothesis can (if true) provide a perfectly satisfactory explanation of this evidence *even if it fails to explain, of a single particular non-black thing, why that non-black thing is a non-raven*. Thus is symmetry restored. So the explanationist resolution of the paradox considered above won't quite do as it stands.

However, help is on the way. Consider again what we want to explain. On the one hand, we have our evidence:

E: A great many ravens have been sought out, in a variety of different settings and conditions, and all ravens observed to date have been found to be black.

then we have our alternative evidence:

E': A great many non-black things have been sought out, in a variety of settings and conditions, and all non-black things observed to date have been found to be non-ravens.

E would be very nicely explained by hypothesis H:

H: All ravens are black.

But then again, it would be explained about as well by this hypothesis:

H*: Approximately all ravens are black.

After all, if non-black ravens are out there, but exceedingly scarce, then this surely accounts for our failure to find one despite our diligent search. Similarly, E' would be explained very nicely by this hypothesis:

H': All non-black things are non-ravens.

But then again, it would be explained just about as well by this one:

H'*: Approximately all non-black things are non-ravens.

So we can see that what E inductively supports, to any very high degree, is not H, but the disjunction of H and H* – or equivalently, H*. And what E' inductively supports, to any very high degree at all, is not H', but the disjunction of H' and H'* – or equivalently, H'*. And though H and H' are logically equivalent, H* and H'* are not.

The apparent paradox was that it looked as though E' lent significant inductive support to a hypothesis that it shouldn't be able to support, namely H; pre-theoretically, it seems that you need to do a diligent search for ravens in order to acquire strong inductive support for a hypothesis about the color distribution of ravens, but at the same time the clever argument made it seem as if we can get support for a hypothesis like that while doing only a very different kind of search. We can now see that the paradoxical appearance was a mere appearance. E' does not lend inductive support to H' after all. The most that it supports is H'*, which is logically independent of H and H*.

## 9. Another Paradox: The Ravens' Revenge?

The Reichenbachian resolutions of the paradox presented above nowhere make the assumption that we are confident that there are more non-black things than there are ravens; moreover, they do not appear to depend on that assumption in any way. This appears to be an important advantage over the standard Bayesian resolution (which they share with many other attempted resolutions of the paradox – see above, section 3). However, things are not quite as they appear. For there is a second paradox, not identical to Hempel's original one but very similar to it, which neither the Reichenbachian resolutions, nor any other resolution, will be able to resolve without making the same assumption. The key to seeing this is found in some of Reichenbach's remarks, quoted and discussed above in section 4.

The key to Reichenbach's resolution of the paradox lies in two steps: first, understanding instantial confirmation not as a way of confirming hypotheses of the form

$\forall x(Ax \supset Bx)$

but instead as a way of confirming hypotheses of the form

(158) $P(Bx|Ax) \approx 1$

with the probability function P interpreted as referring to the actual relative frequencies; second, recognizing that (158) is logically independent of (159):

(159) $P({\sim}Ax|{\sim}Bx) \approx 1$

so that confirmation of the second doesn't automatically make for confirmation of the first. But Reichenbach did not stop with this observation; he also, in effect, determined a special condition under which confirmation of the second would automatically provide confirmation of the first. For he showed that the condition:

(163) $P(Ax) + P(Bx) \geq 1$

is a sufficient condition for the truth of:

(*) $P(Bx|Ax) \geq P({\sim}Ax|{\sim}Bx)$

It follows that if (163) is true, then for any tolerance for error d:

(**) if $(1 - P(\sim Ax|\sim Bx)) \leq d$, then $(1 - P(Bx|Ax)) \leq d$

so for any fixed standard of approximation, if $P(\sim Ax|\sim Bx) \approx 1$, then $P(Bx|Ax) \approx 1$. Reichenbach concludes that when (163) holds, confirming evidence for (159) is automatically confirming evidence for (158).

The probability function P that figures in the above conditions is not the Bayesian's personal probability or credence function; it represents the actual frequencies. Nothing has been said so far about credences. So how can we be sure that Reichenbach is right to say that under the condition (163), confirming evidence for (159) is automatically confirming evidence for (158)?

Suppose our subject assigns credence 1 to (163). Then she must also assign credence 1 to all of the logical and mathematical consequences of (163), including (**). And it follows from that that she must assign at least as high a credence to (158) as she does to (159). It doesn't follow from this that anything that is, for a Bayesian, confirming evidence for (159) is automatically confirming evidence for (158); it is possible that she encounters a piece of evidence E such that conditionalizing on E makes her credence in (159) go up and makes her credence in (158) go down, so long as her credence in (158) ends up at least as high as her credence in (159) does. So while it isn't necessarily the case that every piece of evidence for (159) is going to be a piece of evidence for (158), it is true that our subject's total body of evidence is such that her current credence in (158) is at least as great as her credence in (159). So for example, if her total body of evidence comes to include a great deal of evidence of the sort that supports (159) – for example, if she samples from the non-black things, randomly and under a wide variety of circumstances, and all of the observed non-black things have turned out to be non-ravens, and she finds no other sorts of evidence that tend to depress her credence in (159) – then that same total body of evidence, once she finishes conditionalizing on it, will have raised her credence in (158) at least as high as it has raised (159). It is in this sense that Reichenbach effectively shows that when (163) is true, confirming evidence for (159) is confirming evidence for (158).

This is not the same paradoxical conclusion that Hempel reached in his famous examination of the ravens, but it is a paradoxical conclusion in the same ballpark. One important difference is that it is limited to the case where our subject's credence in (163) is 1. Suppose we are in that case: Then, should our subject come across evidence of the sort we would think of as providing excellent inductive support to (159) – many and varied observed non-black things, all of them non-ravens, for example – then by the time she finishes taking account of that evidence, she will end up very confident in (158), i.e. that about all ravens are black. This offends our basic pre-theoretic judgments about evidence and confirmation. In fact, it offends the very same basic pre-theoretic judgment that Hempel's original paradox offends, namely that if you want to find evidence that supports a general hypothesis about ravens, it's important to observe some ravens; observing green beer bottles, white shoes and so forth simply will not do the trick. Insofar as we embrace this thought, we should find this new paradoxical conclusion just as offensive as we find Hempel's original one.

So what we see here is that, in the case where our subject is certain that $P(A) + P(B) \geq 1$, this basic intuitive judgment of ours is simply false. There is no wiggle room here: Reichenbach's key result here is not a consequence of any philosophical theory he endorses; it's a theorem of the probability calculus – namely, that if $P(Ax) + P(Bx) \geq 1$, then $P(Bx|Ax) \geq P({\sim}Ax|{\sim}Bx)$; hence, if $P({\sim}Ax|{\sim}Bx) \approx 1$, then $P(Bx|Ax) \approx 1$, for any fixed standard of approximation. To show that the offended intuition really is false in this case, we don't need to make any contentious assumptions about confirmaton or evidence at all; we only need to assume that consistency requires a subject who is completely confident of one proposition to be equally confident of the logical and mathematical consequences of that proposition. It follows immediately that for any consistent subject, if $Cr(P(Ax) + P(Bx) \geq 1) = 1$, then:

$$Cr(P(Bx|Ax) \approx 1 \mid E) \geq Cr(P({\sim}Ax|{\sim}Bx) \approx 1 \mid E)$$

no matter what the evidence E is, from which our paradoxical result follows.

Notice that the condition (163), $P(Ax) + P(Bx) \geq 1$, is logically equivalent to:

$$P(Ax) \geq P({\sim}Bx)$$

Since P represents actual frequencies, this means that there are at least as many ravens as there are non-black things. So what we have seen is that if our subject is fully confident that there are at least as many ravens as non-black things, then a conclusion follows which appears paradoxical for the very same reason that the original Paradox of the Ravens seems paradoxical. In reaching this conclusion, we invoked no controversial premises. So no matter which resolution of the Ravens Paradox you take to be the correct one, consistency demands that you too recognize that the same basic pre-theoretic judgments are offended whenever a subject is fully certain that there are at least as many ravens as non-black things. Those resolutions of the Ravens Paradox which, unlike the standard Bayesian one (but like those of Scheffler, Goodman, Armstrong, Lange, and Lipton, as well as the Reichenbachian resolutions presented above), seek to save the offended intuitions rather than explain away their attractiveness to us, cannot ultimately succeed: They too must admit that the offended intuition is simply false in some cases.

What are those cases? What we have seen so far is just that the offended intuitive judgment is false in a case where the subject is fully confident that there are at least as many ravens as non-black things. But a similar result follows from weaker assumptions. Suppose instead that our subject assigns some non-negligible credence c to the condition that there are at least as many ravens as non-black things. (Notice that this is equivalent to supposing that our subject is not extremely confident that there are more non-black things than ravens – that is, that the crucial assumption of the standard Bayesian resolution of the paradox is false.) Then on pain of incoherence she must assign a credence of at least c to every mathematical consequence of this condition. So, she must assign a credence of at least c to the conditional:

$$\text{if } P(\sim Ax|\sim Bx) \approx 1, \text{ then } P(Bx|Ax) \approx 1$$

It also follows that a lower bound on her credence for $P(Bx|Ax) \approx 1$ is her credence of the conjunction of $P(\sim Ax|\sim Bx) \approx 1$ and the proposition that there are at least as many ravens as non-black things:

$$Cr(P(Bx|Ax) \approx 1) \geq Cr(P(\sim Ax|\sim Bx) \approx 1 \ \& \ P(Ax) \geq P(\sim Bx))$$

Now suppose that she discovers evidence of the kind described above – evidence of the sort that seems tailor-made to provide strong inductive support for the hypothesis $P(\sim Ax|\sim Bx) \approx 1$. This evidence might well be evidentially irrelevant to $P(Ax) \geq P(\sim Bx)$; suppose that it is. Then when she conditionalizes on this evidence, the credence on the right-hand side of the last inequality should go up. Thus, a lower bound on her credence in $P(Bx|Ax) \approx 1$ also goes up. It doesn't necessarily follow that her credence in $P(Bx|Ax) \approx 1$ will go up; that depends on the rest of the details about her credence function. But this is enough to show that it can happen that a rational subject who meets our condition (namely, that she has non-negligble credence in the condition that there are at least as many ravens as non-black things) can have her credence in "About all ravens are black" dramatically raised by the sort of the evidence that we would call strong inductive evidence for "About all non-black things are non-ravens" – lots of observations of randomly elected non-black things turning out to be non-ravens, and so forth. Our widely-shared pre-theoretic judgment is that this sort of thing should be impossible. So in the case of a subject who satisfies our condition, that pre-theoretic judgment is simply wrong.

Recall from section 3 that Scheffler advanced the plausible argument that since there are cases in which the original paradox of confirmation seems just as paradoxical as in the ravens case, but in which the assumption made by the standard Bayesian resolution – namely, that the subject starts out very confident that there are more non-black things than ravens – is false, we should prefer a resolution which (unlike the standard Bayesian resolution) works even when that assumption is false. Now we are in a position to see what was wrong with that argument. Even if some other resolution of the original Ravens Paradox is correct, in any case where the standard Bayesian's assumption is false, we have a proof of a conclusion that runs against the very same intuition that the conclusion of the original Ravens Paradox runs against. So in cases where that assumption is false, it is no good trying to save that intuition: It is not salvageable in such cases. It is true that the standard Bayesian resolution cannot tell us why this false intuitive judgment seems so clearly true to us even in those cases. But that is not the standard Bayesian's problem alone; on everyone's view, the intuitive judgment is false in those cases, so we are all equally in need of a way of explaining that intuition away. The apparent advantage that the alternatives to the standard Bayesian resolution seemed to enjoy has thus evaporated.

## 10. Conclusions

Reichenbach's neglected remarks on the Ravens Paradox have yielded interesting fruit of two kinds. First, they indicate a new strategy for resolving the Ravens Paradox, interestingly different from the more familiar resolutions and plausible enough to be worth taking seriously. Second, they point the way to a refutation of an alleged advantage that some non-Bayesian resolution have over the standard Bayesian resolution.[22]

## REFERENCES:


Armstrong, D. M. (1983): *What is a Law of Nature?* Cambridge: Cambridge University Press.

Fitelson, Branden (2006): "The Paradox of Confirmation," *Philosophy Compass* 2006: 95-113.

Fitelson, Branden and Hawthorne, James (2006): "How Bayesian Confirmation Theory Handles the Paradox of the Ravens," in *Probability in Science*, ed. Ellery Eells and James Fetzer, Chicago: Open Court.

Good I. J. (1967): "The White Shoe is a Red Herring," *The British Journal for the Philosophy of Science* 17: 322.

Good, I. J. (1968): "The White Shoe Qua Red Herring is Pink," *The British Journal for the Philosophy of Science* 19: 156-157.


---


[22] Marc Lange and Matt Kotzen gave me enormously helpful comments on an earlier draft.

Goodman, Nelson (1945): *Fact, Fiction and Forecast*, Cambridge, Mass: Harvard University Press.

Hempel, Carl G. (1945): "Studies in the Logic of Confirmation," *Mind* 54: 1-26.

Hosiasson-Lindenbaum, Janina (1940): "On Confirmation," *Journal of Symbolic Logic* 5: 133-148.

Lange, Marc (2000): *Naturl Laws in Scientific Practice*. Oxford: Oxford University Press.

Lipton, Peter (2004): *Inference to the Best Explanation*, Second Edition. London: Routledge.

Lipton, Peter (2007): "The Ravens Revisited," *Philosophy* 61(Supplement): 75-95.

Mackie, J. L. (1963): "The Paradox of Confirmation," *The British Journal for the Philosophy of Science* 52: 265-277.

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

Reichenbach, Hans (1976): *Laws, Modalities, and Counterfactuals*. Berkeley: University of California Press.

Scheffler, Israel (1981): *The Anatomy of Inquiry*. Indianapolis: Hackett.

Vranas, Peter B. M. (2004): "Hempel's Raven Paradox: A Lacuna in the Standard Bayesian Solution," *British Journal for the Philosophy of Science* 55: 545-560.

White, Roger (2005): "Explanation as a Guide to Induction," *Philosopher's Imprint* 5(2): 1-29.