Rerum Causae, 6 (2014), 55-74

Dynamic Epistemic Confirmation: A Qualitative Solution to the Paradox of the Ravens

Chloé de Canson*

Abstract

Fitelson and Hawthorne (2006) argue that Hempel's qualitative response to the paradox of the ravens cannot be satisfactory: they claim firstly that classical deductive logic cannot capture an important distinction needed for Hempel's solution of the paradox, and secondly, that Hempel's solution is internally inconsistent. I show that both of these claims rest on a restrictive understanding of Hempel's solution, and I argue that if confirmation is understood as based on epistemic states, then both these criticisms can be avoided. I finish by giving some reasons to hope that qualitative (dynamic epistemic) and Bayesian confirmation can be unified.

Introduction

The aim of this paper is to argue in favour of a qualitative account of confirmation. More precisely, the paper seeks to show that a dynamic epistemic account of confirmation can be used to support Carl G. Hempel's[1] solution to

* I would like to thank Alexandru Marcoci for his immensely inspiring and helpful comments.
[1] Hempel, 'Studies in the Logic of Confirmation I'.

the paradox of the ravens against Branden Fitelson and James Hawthorne's[2] recent critique, which seeks to show the mismatch between Hempel's theory and his solution to the paradox.

In the first part of the paper, I argue that, in order to account for scientific practice, a theory of confirmation should incorporate both a qualitative and a quantitative dimension. In the second part of the paper, I outline the main elements of Hempel's theory and show how to derive from them the problematic conclusion that the observation of an object like a green teapot confirms the hypothesis that all ravens are black. This is commonly called the paradox of the ravens. I then explain how Hempel deals with this paradox: by upholding that the conclusion, which seems paradoxical, is in fact true. In the third part of the paper, I explain how Fitelson and Hawthorne claim to establish the internal inconsistency of that solution, and show how they have failed to do so, on two different levels. Firstly, they claim that Hempel's theory of confirmation, which is expressed in classical deductive logic, is not sufficiently rich to capture a distinction needed to solve the paradox of the ravens in a qualitative way. I argue that, in fact, if Hempelian confirmation is understood in a dynamic epistemic manner, it is possible to formalise Hempel's solution in classical deductive logic without failing to capture any of its important aspects. Secondly, Fitelson and Hawthorne argue that, if there were a way to understand Hempel's solution in a classical deductive logic, this would lead to an inconsistency. I show that this is not the case: if, again, confirmation is based on the knowledge states of the agent, then this inconsistency does not arise. I conclude that Fitelson and Hawthorne give us no reason to reject Hempel's solution to the paradox of the ravens. In the fourth part of the paper, I finish by giving some reasons to believe that a unified account of confirmation (integrating both dynamic epistemic confirmation and Bayesian confirmation) is possible.

[2] Fitelson and Hawthorne, 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens'.

I. The Complementarity of Qualitative and Bayesian Confirmation

In this part of this paper, I present the notion of confirmation and give some reasons to uphold a qualitative theory of confirmation, as well as a Bayesian one.

I first present the intuitive notion of confirmation. When we have a hypothesis H and we encounter a positive instance of this hypothesis (evidence E), we want to say that E confirms H. For example, take H to be the hypothesis that all swans are white. The observation of a white swan, E, intuitively confirms the hypothesis H. This means that E supports H, makes it more plausible, more acceptable, more defendable. In Popperian terms, E corroborates H.[3] Intuitively, any hypothesis, as long as it has some sort of empirical consequences, can be confirmed in that sense: if the empirical consequences of a theory are observed, then it is more plausible that the theory is true.

Such confirmation is not only intuitively appealing, but also very important for scientific practice. Indeed, a major criterion for the success of scientific theories is its predictive success. This means that, for a theory to be accepted as successful, it needs for its consequences to have been tested empirically, and for the result of this experiment to be positive. Predictive success and confirmation are intimately linked: if H predicts E (and thus) has predictive success, then E confirms H. This can be illustrated using an example from the history of science. Albert Einstein's General Relativity Theory (GRT) implies that light bends around massive objects. This, in turn, implies that stars should appear further away during the day than during the night. Indeed, during the day, the light from the stars to the observers bends around the sun (see full line below). However, during the night, since the sun is not between the stars and the earth, the light travels in a straight line (see dotted line below).

[3] Popper, 'Science: Conjectures and Refutations'.

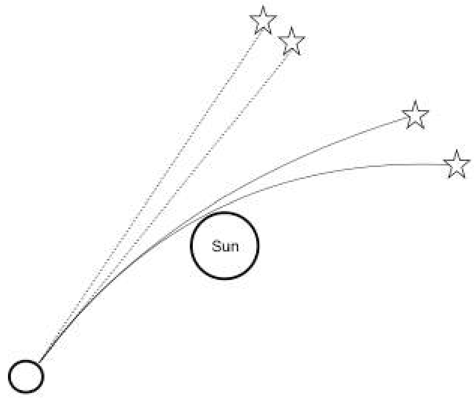

Arthur Eddington, a British astronomer, observed this phenomenon – a consequence of GRT. This observation was taken to be one of the most important confirmations of Einstein's theory, and contributed greatly to its success.

I now give some reasons to uphold both a qualitative and a Bayesian account of confirmation. The Bayesian account of confirmation, which is based on probabilities, has been recently regarded as the best account of confirmation. Indeed, it has many advantages over a qualitative one, and these have been discussed at length in the literature. One of these is that it gives us the tools to determine, when we have two pieces of evidence E1 and E2, which one confirms a hypothesis H the most. For example, both the facts that light bends around massive objects and that small objects dropped near the surface of the earth fall to the ground are predicted by GRT, and therefore confirm it. However, Eddington's test was taken to confirm GRT more than the fall of a pen when it is dropped. Bayesian confirmation accounts for this intuition very well.[4]

[4] Suppose that a theory H has two empirical consequences E_1 and E_2: we have $H \subseteq E_1$ and $H \subseteq E_2$. Now, supose that E_1 is more likely to be true than E_2: $p(E_1) > p(E_2)$. Intuitively, the observation of E_2 confirms H more than the observation of E_1 does. This can be accounted for very well in Bayesian confirmation. Indeed, using Bayes' theorem, we have:

$$p(H|E_1) = \frac{p(H).p(E_1|H)}{p(E_1)}, \text{ given } p(E_1) \neq 0;$$
$$p(H|E_2) = \frac{p(H).p(E_2|H)}{p(E_2)}, \text{ given } p(E_2) \neq 0.$$

However, after Eddington's test results were published, the question in the scientific community was not primarily whether this confirmed GRT *more* than the fall of a pen, or *to what extent* this confirmed GRT, but rather *whether*= it did. The primary question when it comes to confirmation is inherently qualitative. I argue in this paper that Hempel's account of confirmation (understood so that it is concerned with states of knowledge as evidence, rather than objects or propositions) can give us precisely this: a criterion for determining solely whether some piece of evidence confirms a hypothesis. It can be argued that Bayesian confirmation can also give us this result. However, this requires appeal to probability (one would need to calculate the prior probability of H and the posterior probability of H given E), which does not seem to coincide with all the confirmation questions asked in science: when one learns that light bends around the sun and claims that this confirms GRT, that person does not calculate the probabilities of GRT and GRT|E. A good account of whether some evidence confirms a hypothesis therefore should not appeal to probabilities – or at least not in all cases.

Therefore, a theory of confirmation which would be composed of both a qualitative (Hempelian) and a quantitative (Bayesian) component would therefore be much more complete a theory that if it were concerned only with one of these two aspects. Indeed, such a theory would enable the scientist to determine whether E confirms H, as well as which of E and E' confirms H more and (potentially) to what extent E confirms H.

Since $H \subseteq E_1$ and $H \subseteq E_2$, $p(E_1|H) = p(E_2|H) = 1$. Therefore:

$$p(H|E_1) = \frac{p(H)}{p(E_1)} < \frac{p(H)}{p(E_2)} = p(H|E_2)$$

.

II. Hempel's Qualitative Theory of Confirmation and his Solution to the Paradox of the Ravens

In this second part of the paper, I first outline the basic idea and some elements of Hempel's qualitative theory of confirmation. I then show how the paradox of the ravens can be derived from the components of that theory, and explain how Hempel responds to this paradox.

Hempel made explicit the characteristics that a satisfactory theory of qualitative confirmation should have. According to him, the relation of confirmation should be a *logical* relation (which I henceforth formalise with a 'Hempelian turnstile' $\vdash_H$), which should not involve any numerical values, and which should satisfy certain conditions.[5] I outline the two such conditions which are relevant for the paradox of the ravens.

The first one is the *equivalence condition*.[6] It states that if some evidence E confirms a hypothesis H, and if H is logically equivalent (according to the rules of classical logic) to another hypothesis H', then E confirms H'. This fits our intuition. Indeed, if two hypotheses H and H' are logically equivalent, then they are simply two different formulations of the same idea. Confirmation of a hypothesis should not depend on the terminology, and therefore E should confirm both H and H'. This equivalence condition (EC) can be expressed formally as follows, where $\equiv$ is logical equivalence:

(EC) If $E \vdash_H H$ and $H \equiv H'$
Then, $E \vdash_H H'$

We have seen that the theory of qualitative confirmation's main object of enquiry is the relation between hypothesis and evidence. In the most simple confirmation cases, E is a positive instance of H. For example, E 'a is a black raven' is a *positive instance* of H 'all ravens are black', and therefore confirms

[5] Hempel, 'Studies in the Logic of Confirmation I' 1-3.
[6] Hempel 'Studies in the Logic of Confirmation II' 103-105.

it. Patrick Maher points out that Hempel implied that this confirmation is relative to 'tautological background' corpus.[7] Indeed, in Hempel, the question is whether learning that E *and nothing else* confirms H. This is called the Nicod condition (NC). The above example can be formalised in this way:

(NC) $(Ra \wedge Ba) \vdash_H \forall x(Rx \rightarrow Bx)$, given tautological background knowledge.

The Hempel paradox, or paradox of the ravens, can be derived from Hempel's theory. This is done in the following way. Let us take H to be the hypothesis 'all ravens are black' $[\forall x(Rx \rightarrow Bx)]$. Then, let H' be 'all non-black things are non-ravens' $[\forall x(\neg Bx \rightarrow \neg Rx)]$. By contraposition, H and H' are equivalent. The evidence E 'a is a non-black non-raven' $[(\neg Ra \wedge \neg Ba)]$, such as a green teapot, is a positive instance of H' and therefore, by the Nicod condition of confirmation (NC), confirms H'. Furthermore, since H $\equiv$ H', by the equivalence condition of confirmation (EC), E confirms H. This means that the observation of a green teapot confirms the hypothesis that all ravens are black. Nelson Goodman commented on this conclusion: the theory of qualitative confirmation gives puzzling 'prospects for indoor ornithology.'[8] Indeed, intuitively, the conclusion seems flagrantly false. This appears to be an important problem for the theory of qualitative confirmation: since the paradoxical conclusion is logically derived from it, if the conclusion indeed is false, then the theory is internally inconsistent.

In order to resolve the paradox of the ravens, we have four possible options. We could reject the Nicod condition, or the equivalence condition, or the claim that H and H' are equivalent (that is, reject at least one rule of classical logic), or the claim that the conclusion is paradoxical. Denying the (NC) or the (EC) would make the theory of qualitative confirmation counterintuitive: I

[7] Maher, 'Inductive Logic and the Ravens Paradox'.

[8] Goodman, *Fact, Fiction and Forecast* 69.

have explained above how these conditions were intuitively good conditions for qualitative confirmation. Rejecting contraposition would mean rejecting classical logic as the logic in which to study confirmation. I will argue in the next section that classical deductive logic is in fact very well fitted to the study of confirmation. Thus, the last possibility to resolve the paradox is to claim that the conclusion, although it seems false, is in fact true: the fact that I can observe a green teapot does confirm the hypothesis that all ravens are black.

And this is what Hempel did: he asserted that the paradoxical conclusion (PC: if one observes that an object about which nothing is antecedently known[9] is a non-black non-raven, then this observation confirms the hypothesis that all ravens are black) is only *apparently paradoxical*. The paradox of the ravens, according to Hempel, is in fact not a paradox, but simply a problematic case. He argues that the reason that one might be misled into thinking otherwise is by conflating (PC) with a different claim (PC*: if one observes that an object which is already known to be a non-raven is non-black, then this observation confirms the hypothesis that all ravens are black).[10] (PC*) is obviously false. Indeed, if we already know that the object we observe is not a raven, it can be neither a confirming instance nor a counterexample to the hypothesis H. Observing a known non-raven cannot tell us anything about the colour of ravens. On the other hand, (PC) is true. Let us assume that there is a finite number of macroscopic objects in the world. There is a certain number of them, call it n, which have not yet been observed. There is an object in a room about which we know nothing. If we observe that object to be a non-black non-raven, then

[9] It might seem that the condition of nothing being antecedently known springs out of nowhere. However, we had defined the Nicod condition as being that a positive instance of a theory confirms that theory, *relative to no or tautological background knowledge*, in accordance with what Maher showed Hempel to have in mind. Therefore, according to this (NC), what confirms that all ravens are black is the observation of a green teapot on its own – without any other (non-tautological) background knowledge: nothing is antecedently known about the green teapot before its observation.

[10] Fitelson, 'The Paradox of Confirmation' 65.

the number of possible objects which would falsify the hypothesis (that is, the number of possible non-black ravens) decreases from n to n-1. In that sense, because there are fewer possibilities for H to be false, E' confirms H. Thus, (PC) is true, and the paradox of the raven is only an apparent one.[11]

By clarifying what is actually contained in the conclusion of the paradox of the ravens, Hempel has shown that this conclusion is in fact intuitively true, and has thus dismissed the paradox of the ravens as non-paradoxical. This solution seems to be perfectly acceptable. It has however been criticised, notably by Fitelson and Hawthorne who have claimed that this solution is internally inconsistent. In the third part of this paper, I evaluate and reject this criticism.

III. Fitelson and Hawthorne's Criticism of Qualitative Confirmation and Dynamic Epistemic Confirmation

Fitelson and Hawthorne argue that Hempel's qualitative response to the paradox of the ravens is not satisfactory: they claim firstly that classical deductive logic cannot capture an important distinction needed for Hempel's solution of

[11] This might seem like a Bayesian answer to the paradox, because 'there are fewer possibilities for H to be false' may be understood as 'the probability of H being true increases'. In case the reader is dissatisfied with this solution, Fitelson and Hawthorne (2006, p. 8) give us another way to resolve this paradox in a qualitative way. According to them, a Hempelian story can be told along these lines. H implies, given that the object is not black ($\neg Ba$), that it is not a raven ($\neg Ra$). This means that, given $\neg Ba$, Ra falsifies H, and so $\neg Ra$ confirms H. Since H entails neither Ra nor $\neg Ra$ given $\neg Ba$, we can say that given $\neg Ba$, whether a is a raven is neutral with respect to H (it neither confirms nor falsify it). So, overall, given Ba or $\neg Ba$, and given that we have no information on whether a is a raven, $\neg Ra$ confirms H. This is a justification of (PC). (PC*) on the other hand, can be said to be false for the following reason. Given $\neg Ra$, H entails neither Ba nor $\neg Ba$. So, given $\neg Ra$, neither Ba nor $\neg Ba$ can confirm H. Thus, (PC*) is false.

the paradox, and secondly, that Hempel's solution is internally inconsistent.[12] I examine these criticisms in turn and show that they do not threaten Hempel's solution. Indeed, if confirmation is understood as based on the states of knowledge of an agent, then both of these criticisms can be avoided: dynamic epistemic confirmation is a form of qualitative confirmation which escapes Fitelson and Hawthorne's objections.

a. The Paradox of the Ravens and Classical Deductive Logic

Fitelson and Hawthorne note that there is no distinction in classical deductive logic between the following two kinds of claims:

- $\neg Ba$ confirms $\forall x(Rx \to Bx)$ given $\neg Ra$
- $\neg Ba \wedge \neg Ra$ confirms $\forall x(Rx \to Bx)$

Furthermore, they remark that the distinction between 'and' and 'given that' seems to be the crucial difference between (PC) and (PC*). Indeed, in (PC), the confirmation is done by the non-ravenness and the non-blackness of the green teapot, taken together. However, in (PC*), the confirmation (or rather lack of confirmation, since the claim is false) is done by the non-blackness of the green teapot, given that it is not a raven. The non-ravenness and the non-blackness of the teapot are separated. Fitelson and Hawthorne conclude that, since the distinction between 'and' and 'given that' is the crucial difference between (PC) and (PC*); and since classical deductive logic, the framework in which Hempel's theory is set out, does not capture this distinction, Hempel's solution cannot be understood in Hempel's theoretical framework. Therefore, they come to the conclusion that Hempel's qualitative solution to the paradox of the ravens is not satisfactory, and instead argue in favour of a Bayesian solution

[12] Fitelson and Hawthorne, 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens'.

to the paradox.[13]

However, I argue that there is a way to express (PC) and (PC*) in classical deductive logic in such a way that the distinction between the two is clear. I also argue that both Hempel and Fitelson and Hawthorne hint at this account.

The way in which to formalise (PC) and (PC*) is dependent upon the answer to the question: what entity does the confirmation? Fitelson and Hawthorne claim that Hempel slides back and forth between 'objectual and propositional senses of confirmation.'[14] They seem to consider that, in Hempel's theory, either the objects or the propositions do the confirming. That is, either the green teapot confirms the hypothesis that all ravens are black, or the claim '*a* is a green teapot' does. However, there is a third possible way to understand confirmation: it is the *knowledge* that *a* is a green teapot which confirms the hypothesis that all ravens are black. Fitelson and Hawthorne, although they do not acknowledge this understanding of confirmation, hint at it as they emphasise the importance of the 'confirmational impact of *learning* E', and as they claim that 'this *observation* [the observation of a non-black non-raven – and therefore the coming to know that the object is a non-black non-raven] confirms that all ravens are black' (my italics in the two quotes).[15] Hempel, too, suggests that knowledge is highly relevant to confirmation. He writes: 'if we are careful to avoid this[16] tacit

[13]The Bayesian account of confirmation, because it is based on probabilities, is perfectly suited for capturing the distinction between 'and' and 'given that'. Indeed, the claim that $\neg Ba \wedge \neg Ra$ confirms $\forall x(Rx \rightarrow Bx)$ is understood in a Bayesian way as $p(\forall x(Rx \rightarrow Bx) \mid \neg Ba \wedge \neg Ra) > p(\forall x(Rx \rightarrow Bx))$; whereas $\neg Ba$ confirms $\forall x(Rx \rightarrow Bx)$ given $\neg Ra$ can be formalised as $p(\forall x(Rx \rightarrow Bx) \mid \neg Ba \wedge \neg Ra) > p(\forall x(Rx \rightarrow Bx) \mid \neg Ra)$. The two are clearly distinct.

[14]Fitelson and Hawthorne, 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens' 9.

[15]Fitelson and Hawthorne 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens' 6-7.

[16]Hempel writes this in the context of his example of the sodium salt. However, this is completely relevant for the paradox of the ravens: the problem in both of these examples is the same. Therefore, if we transpose this sentence to the context of the raven's paradox, 'this tacit reference to additional knowledge' refers to the assumption that the object observed is known

reference to additional knowledge, [...] the paradox vanishes' (my italics).[17] This seems to suggest that it is the *discovery* that a is a green teapot which confirms that all ravens are black, not the *fact* that a is a green teapot. And indeed, this is very plausible: we say that a theory has been confirmed when we learn the evidence that supports it. Let us illustrate this using Eddington's confirmation of GRT. If confirmation were understood 'objectually', then the bending of the light around the sun would be what confirms Einstein's theory. But since the light has always bent around the sun, the theory would always have been confirmed, and Eddington's experiment would not have had any confirmational impact with respect to GRT. This is clearly counterintuitive, as what is taken by scientists to 'do the confirming' is Eddington's experimental results. This means that if we want to have a plausible theory of confirmation that fits with our intuitions, we need to uphold that the discovery, or more generally the evolution of knowledge, of the light bending is what confirmed GRT – that is, confirmation must be understood as based on knowledge revision. I call this the *dynamic epistemic* sense of confirmation, as opposed to the objectual and propositional senses proposed by Fitelson and Hawthorne.

Now that we have established that confirmation should be understood in a dynamic epistemic manner, let us come back to the problem of whether formalising (PC) and (PC*) in classical deductive logic is possible. Fitelson and Hawthorne had claimed that it was not because, in classical deductive logic, there is no distinction between '$\neg Ba$ confirms $\forall x(Rx \rightarrow Bx)$ given $\neg Ra$' and '$\neg Ba \wedge \neg Ra$ confirms $\forall x(Rx \rightarrow Bx)$'. This is indeed the case. However, since it is not the object or proposition $\neg Ra$ which does the confirming, this remark is no longer relevant to the formalisation of (PC) and (PC*). Instead, let us use a knowledge operator K, accompanied by time subscripts. So, let 'it is known at time t that a is not a raven' be $K_t(\neg Ra)$. Using this knowledge operator makes it possible to formalise, in a classical deductive logic, (PC*) and (PC) such that

not to be a raven.

[17] Hempel, 'Studies in the Logic of Confirmation II' 20.

the difference between the two is evident. Indeed, it suffices to express that, at time t, in (PC*) there is $K_t(\neg Ra)$, and in (PC) there is not; but in both there is $K_{t+1}(\neg Ra) \wedge K_{t+1}(\neg Ba)$. This captures the intuition that observing at once an object which is non-black and non-raven is different from observing that object which was previously known not to be a raven is not black.

Therefore, I have argued that, contrary to what Fitelson and Hawthorne claim, there is a way to differentiate (PC) and (PC*) in classical deductive logic. This is dependent on understanding the evidence as being, not an object or a proposition, but a piece of knowledge – which the papers of both Hempel and Fitelson and Hawthorne suggest. Thus dynamic epistemic confirmation, since it is based on classical deductive logic and permits a differentiation between (PC) and (PC*), evades Fitelson and Hawthorne's first criticism of qualitative confirmation.

b. The Paradox of the Ravens and the Internal Inconsistency of Hempel's Solution

Fitelson and Hawthorne point out that Hempel, in his theory of qualitative confirmation, is committed to monotonicity (M), which is the claim that if some E confirms H, then E $\wedge$ Γ also confirms H, for any Γ.[18] Therefore, in their words, 'if (PC) is true, then (PC*) must be true'.[19] They thus conclude that Hempel's evasion of the paradoxical conclusion fails. This argument can be attacked on two levels: the first is by arguing that monotonicity does not hold for relations of confirmation, and the second is by arguing that monotonicity cannot permit us to derive the truth of (PC*) from the truth of (PC). I begin this section by discussing briefly whether monotonicity does in fact hold – I argue that is does not in general, but that in one case relevant to the present paradox, it might do.

[18] Fitelson and Hawthorne, 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens' 8.

[19] Fitelson and Hawthorne, 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens' 8-9.

I then show that, even if monotonicity of confirmation is granted, it cannot be used to derive the truth of (PC*) from the truth of (PC). I conclude that Fitelson and Hawthorne's argument against Hempel's solution to the paradox is not tenable.

I begin by questioning whether confirmation is monotonic. It is obvious that, in many cases, it is not. For example, take E_1 to be '*a* is a black raven' and E_2 to be '*a* is a white raven'. E_1 clearly confirms the hypothesis that all ravens are black, but $E_1 \wedge E_2$ certainly does not – in fact, it falsifies it. However, a weakened version of (M) can maybe be granted: if $E_1 \vdash_H H$ and it is not the case that E_2 either disconfirms or falsifies H (that is, E_2 either confirms or is neutral with respect to H), then $E_1 \wedge E_2 \vdash_H H$. This is rather plausible: if E_1 is the result of Eddington's test and E_2 is the knowledge that it rained yesterday, although E_2 is irrelevant to GRT, the conjunction of both pieces of evidence does confirm the theory. Let us henceforth take (M) to be this weakened version of monotonicity.

I now show that Fitelson and Hawthorne's claim that 'if (PC) is true, then (PC*) must be true' is mistaken on the dynamic epistemic understanding of confirmation. I begin by formalising (PC*) and (PC) using the knowledge operator described above. I then show that there is no Γ such that (PC) $\wedge$ Γ is equivalent to (PC*). Finally, I explain why Fitelson and Hawthorne might have thought their claim to be true: because of a mistaken interpretation of (PC).

Let us first formalise (PC*). It claims that if it is known at time t that *a* is a non-raven [$K_t(\neg Ra)$] (and therefore it is not the case that at time t, it is known that *a* is a raven [$\neg K_t(Ra)$]); and then at time t+1 it is observed that *a* is a non-black [$K_{t+1}(\neg Ba)$] non-raven [$K_{t+1}(\neg Ra)$], then this observation confirms the hypothesis that all ravens are black [$\forall x(Rx \to Bx)$]. So, its formalisation using our language is:

(PC*): $[K_t(\neg Ra) \wedge \neg K_t(Ra) \wedge K_{t+1}(\neg Ra) \wedge K_{t+1}(\neg Ba)] \vdash_H \forall x(Rx \rightarrow Bx)$[20]

We now come to (PC). There are two possible ways to understand this claim. The first one is to consider that it claims only that it is known at t+1 that *a* is a non-black $[K_{t+1}(\neg Ba)]$ non-raven $[K_{t+1}(\neg Ra)]$. In this interpretation of (PC), there is a mere absence of knowledge at time t – and so we call this first version (PC_{mere}). In formal terms, we have:

(PC_{mere}): $[K_{t+1}(\neg Ra) \wedge K_{t+1}(\neg Ba)] \vdash_H \forall x(Rx \rightarrow Bx)$

However, this is not a satisfactory interpretation of (PC). Indeed, as Fitelson and Hawthorne themselves note, what is important is ‘the confirmational impact of *learning* E’[21] (my italics). This implies that E was not known before its observation. This prior lack of knowledge is necessary to the confirmation, and so, should be made explicit. We therefore make another version of (PC), which I call (PC_{exp}), and which claims that if, at time t, one does not know whether an object *a* is a raven $[\neg K_t(\neg Ra) \wedge \neg K_t(Ra)]$, and if at moment t+1 that same person discovers that *a* is a non-black $[K_{t+1}(\neg Ba)]$ non-raven $[K_{t+1}(\neg Ra)]$, then this discovery confirms the hypothesis that all ravens are black $[\forall x(Rx \rightarrow Bx)]$. ($\text{PC}_{exp}$) can be expressed formally:

(PC_{exp}): $[\neg K_t(\neg Ra) \wedge \neg K_t(Ra) \wedge K_{t+1}(\neg Ra) \wedge K_{t+1}(\neg Ba)] \vdash_H \forall x(Rx \rightarrow Bx)$

I henceforth take (PC_{exp}) to be the correct formalisation of (PC). I have thus formalised (PC*) and (PC) in classical deductive logic, using the dynamic epistemic understanding of confirmation.

I now show, using these formalisations, that it is not possible for (PC*) to follow from (PC) and (M). For this to be true, there would need to exist

[20] Here, the second conjunt of the antecedent $(\neg K_t(Ra))$ is entailed by the first $(K_t(\neg Ra))$. The reason for which I make it explicit in the antecedent is in order to make clear that the only distinction between (PC*) and (PC_{exp}) – see below – is the state of the agent’s knowledge in $\neg Ra$ at time t.

[21] Fitelson and Hawthorne, ‘How Bayesian Confirmation Theory Handles the Paradox of the Ravens’ 7.

a Γ which, conjoined with $\neg K_t(\neg Ra) \wedge \neg K_t(Ra) \wedge K_{t+1}(\neg Ra) \wedge K_{t+1}(\neg Ba)$, is equivalent to $K_t(\neg Ra) \wedge \neg K_t(Ra) \wedge K_{t+1}(\neg Ra) \wedge K_{t+1}(\neg Ba)$. Because the last three conjuncts in the antecedents of (PC) and (PC*) are the same, let us simplify the problem and show that there is no Γ such that $[\Gamma \wedge \neg K_t(\neg Ra)] \equiv K_t(\neg Ra)$. In any classical logic, there does not exist such a Γ unless both Γ and $K_t(\neg Ra)$ are logical falsehoods, which is not the case here.[22] Indeed, the term on the right-hand-side is the negation of one of the conjuncts on the left-hand-side. Therefore, for Γ to exist and for the argument to be valid, Γ would need to cancel out the negation in front of $\neg K_t(\neg Ra)$ so as to obtain $K_t(\neg Ra)$. This, at least in classical logic, is impossible. In this light, the simplicity with which Fitelson and Hawthorne dismiss Hempel's solution as being inconsistent is unwarranted: it is not the case that, in classical logic, 'if (PC) is true, then (PC*) must be true'.

It seems to me that the reason why Fitelson and Hawthorne make the mistaken claim that (PC*) follows from (PC) and (M) is that they consider (PC) to be (PC_{mere}) and not (PC_{exp}); that is, they consider that merely knowing the evidence at time t+1, rather than *discovering* it at t+1, is the source of the confirmation. If this were the case, it could be thought that (PC*) follows from (PC_{mere}) and (M): one would simply use (M) to add the two premises $K_t(\neg Ra)$ and $\neg K_t(Ra)$ to (PC_{mere}), and obtain (PC*).

Therefore, I have shown that Fitelson and Hawthorne's claim that 'if (PC) is true, then (PC*) must be true' is mistaken. I formalised (PC*) and (PC) using the dynamic epistemic sense of confirmation. I then showed that there is no Γ such that (PC)$\wedge\Gamma$ is equivalent to (PC*). Finally, I explained why Fitelson and Hawthorne might have thought their claim to be true.

[22] $[\Gamma \wedge \neg A] \equiv A$ is valid if and only if all interpretations making $\Gamma \wedge \neg A$ true also make A true. This is possible only if both Γ and A are logical falsehood. Here, A is $K_t(\neg Ra)$. It is clearly not a logical falsehood that one knows at time t that object *a* is not a raven. Therefore, there does not exist a Γ such that the inference is valid.

Fitelson and Hawthorne argue that Hempel's qualitative response to the paradox of the ravens is not satisfactory: firstly, they claim that classical deductive logic cannot capture an important distinction needed for Hempel's solution of the paradox, and secondly, that Hempel's solution is internally inconsistent. I have, in this third part of the paper, examined these criticisms in turn and shown that they do not threaten Hempel's solution if confirmation is understood as based on the evolution of knowledge in time. Indeed, I began by showing that there is a way to formalise (PC) and (PC*) in classical deductive logic, provided that confirmation is understood not in the objectual or propositional sense, but in the dynamic epistemic sense. Then, I showed that, with this formalisation of (PC) and (PC*), Fitelson and Hawthorne's argument for the internal inconsistency of Hempel's argument is not valid. I conclude that Hempel's solution to the paradox of the ravens, if it is clarified so that it is understood in a dynamic epistemic manner, is satisfactory.

IV. A Hempelian-Bayesian resolution?

In this last part of the paper, I examine briefly whether a Hempelian-Bayesian solution to the paradox of the ravens is possible, by outlining some similarities between the two and arguing that they give us good hope that a unified account is possible.

There are two important similarities between the qualitative solution given above and most Bayesian solutions: the confirmational impact of learning E, and the interpretations of (PC) and (PC*). The first of these similarities is that both emphasise that the confirmation is done by the discovery of some evidence. Indeed, in Bayesian confirmation, some evidence E confirms a hypothesis H if the probability of H given E is higher than the probability of H alone. The learning of E therefore heightens the probability of H being true and therefore confirms H. This is identical to the dynamic epistemic qualitative account of confirmation that I outlined.

The second similarity between the qualitative and the quantitative solutions of the ravens paradox is that they both distinguish between (PC) and (PC*) on the basis of whether $\neg Ra$ is known before the confirmation. Indeed, Fitelson and Hawthorne express the distinction between the two from a Bayesian point of view as:[23]

- (PC): $p(\forall x(Rx \rightarrow Bx) \mid \neg Ba \wedge \neg Ra) > p(\forall x(Rx \rightarrow Bx))$
- (PC*): $p(\forall x(Rx \rightarrow Bx) \mid \neg Ba \wedge \neg Ra) > p(\forall x(Rx \rightarrow Bx) \mid \neg Ra)$

In this formulation, (PC) expresses that our degree of belief in H given that we know $\neg Ba \wedge \neg Ra$ is higher than our degree of belief in H alone – therefore, learning $\neg Ba \wedge \neg Ra$ confirms H. The qualitative (PC) claimed the same thing. The above Bayesian (PC*), which means that our degree of belief in H knowing that *a* is not a raven is greater once we learn that *a* is not black, thus also expresses the same thing as the qualitative (PC*).

Therefore, the two solutions are similar: they understand confirmation as being 'done' by the same entity (the knowledge of the evidence), and they are both based on the same basic understanding of the claims (PC) and (PC*).

To conclude on this last section, both the qualitative account of confirmation developed in the previous section and the Bayesian account are consistent – and complementary. They both agree that some H can be confirmed by learning some evidence E; and their solutions to the paradox of the ravens are essentially the same in that they are based on the difference between (PC) and (PC*) – which both understand in the same way. So, we can say that the basic claims of these two theories are consistent. This gives us good hope that a unified theory of confirmation (which would integrate a more developed theory of dynamic epistemic confirmation and Bayesian confirmation) is possible. This is

[23] Fitelson and Hawthorne, 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens' 10.

particularly desirable as, as I argued in section I, both types of theories capture different aspects of scientific practice: the qualitative account permits us to determine whether some E confirms some H, whereas the Bayesian account enables us to determine which of some E and some E' confirms H more, and even maybe to what extent E confirms H. A theory that would unify these two would therefore account for all dimensions of scientific practice.

Conclusion

In this paper I have sought to argue in favour of a qualitative account of confirmation. More precisely, I have tried to show that a dynamic epistemic account of confirmation can be used to support Carl G. Hempel's solution to the paradox of the ravens against Branden Fitelson and James Hawthorne's recent critique, which seeks to show the mismatch between Hempel's theory and his solution to the paradox. I first gave a general argument in favour of qualitative confirmation, as well as quantitative. Then, I outlined the main elements of Hempel's theory, showed how to derive from them the paradox of the ravens and explained Hempel's solution to this paradox. Subsequently, I explained how Fitelson and Hawthorne criticise this solution: they argue that it cannot be understood in classical deductive logic, and that it is inconsistent with Hempel's theory. I showed that, in fact, if confirmation is understood in a dynamic epistemic sense (which both Hempel and Fitelson and Hawthorne hint at), both of these criticisms can be overcome. I concluded that Fitelson and Hawthorne give us no reason to reject Hempel's solution to the paradox of the ravens. Finally, I gave some reasons to believe that a unified account of confirmation (integrating both dynamic epistemic confirmation and Bayesian confirmation) is possible.

References

[1] Fitelson, B. and J. Hawthorne. 'How Bayesian Confirmation Theory Handles the Paradox of the Ravens'. *Probability in Science*, ed. E. Eells and J. Fetzer. Chicago: Open Court, 2006.

[2] Fitelson, B. and J. Hawthorne. 'The Wason Task(s) and the Paradox of Confirmation'. *Philosophical Perspectives* 2010; 24(1): 207-241.

[3] Fitelson, B. 'The Paradox of Confirmation'. *Philosophy Compass* 2006; 1(1): 95-113.

[4] Goodman, N. *Fact, Fiction and Forecast*. 4th ed. Cambridge, Mass.: Harvard University Press; 1983.

[5] Hempel, C.G. 'Studies in the Logic of Confirmation I'. *Mind* January 1945; 54: 1-26.

[6] Hempel, C.G. 'Studies in the Logic of Confirmation II'. *Mind* April 1945; 54: 97-121.

[7] Maher, P. 'Inductive Logic and the Ravens Paradox', *Philosophy of Science* 1999; 66(1): 50-70.

[8] Popper, K.R. 'Science: Conjectures and Refutations', *Conjectures and Refutations*. London: Routledge; 1969.

Chloé de Canson is a BSc Philosophy, Logic and Scientific Method student at the London School of Economics (2012-2015). Her main field of interest is philosophy of science, and she hopes to continue in research in that field. You can contact her at [c.de-canson@lse.ac.uk].