UNIVERSITY OF LEEDS

This is a repository copy of *What Is Bayesian Confirmation for?*.

White Rose Research Online URL for this paper:
http://eprints.whiterose.ac.uk/128769/

Article:
Bradley, D (2018) What Is Bayesian Confirmation for? International Studies in the Philosophy of Science, 31 (3). pp. 229-241. ISSN 0269-8595

https://doi.org/10.1080/02698595.2018.1463692

© 2018 Open Society Foundation. This is an Accepted Manuscript of an article published by Taylor & Francis in International Studies in the Philosophy of Science on 25/05/2018, available online: http://www.tandfonline.com/10.1080/02698595.2018.1463692

Reuse
Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

White Rose university consortium
Universities of Leeds, Sheffield & York

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

What Is Bayesian Confirmation for?

Darren Bradley

Department of Philosophy, University of Leeds

CONTACT Darren Bradley / d.j.bradley@leeds.ac.uk / Department of Philosophy, University of Leeds, Woodhouse Lane, Leeds LS2 9JT, UK

ABSTRACT

Peter Brössel and Franz Huber in 2015 argued that the Bayesian concept of confirmation had no use. I will argue that it has both the uses they discussed—it can be used for making claims about how worthy of belief various hypotheses are, and it can be used to measure the epistemic value of experiments. Furthermore, it can be useful in explanations. More generally, I will argue that more coarse-grained concepts (like confirmation) can be useful, even when we have more fine-grained concepts (like credences).

1. Introduction

A centrepiece of contemporary Bayesianism is the Bayesian analysis of the concept of confirmation:

> E confirms H relative to background assumptions B and probability function P if and only if $P(H|E\&B) > P(H|B)$[1]

Peter Brössel and Franz Huber raise an important and neglected question: what is the purpose of the Bayesian conception of confirmation? They consider and reject two possible answers:

(1) that the purpose of the Bayesian conception of confirmation is for making claims about how 'worthy of belief various hypotheses are';

(2) that the Bayesian conception of confirmation can be used to measure 'the epistemic value of experimental outcomes' (Brössel and Huber 2015, 737), and thus to decide which experiments to carry out.

I will argue that the Bayesian conception of confirmation can be used for both purposes. The more general moral is that there are two reasons why coarse-grained concepts can be more useful than fine-grained concepts—they are useful when we are ignorant of the details and they are useful when omitting details improves an explanation.

2. Set-up

Let's put the debate into context by asking a more general question: what is the purpose of a conceptual analysis?[2] The purpose we'll focus on is that it solves what Frank Jackson (1998) calls the location problem—that of connecting the vocabulary of one subject-matter with the vocabulary of a different, better understood, subject-matter. Prior to finding an analysis, we might worry that a concept is incoherent or fails to refer to anything in the world. A failure to find a conceptual analysis can motivate eliminativism about the subject-matter. By contrast, a successful conceptual analysis vindicates the subject-matter.[3]

The concept of confirmation seems to be extensively used by scientists, so part of the interest in giving an analysis of the concept of confirmation is to make sense of this confirmation-talk by 'locating' confirmation using better understood concepts. The Bayesian analysis does this—it defines confirmation in terms of concepts Bayesians already endorse i.e. probability functions and background beliefs. So this analysis might be considered a success of Bayesianism.

The challenge posed by Brössel and Huber is that the concept of confirmation offered by the Bayesian is not useful. What would it take for a concept to be useful? 'Useful' is such a vague and flexible word that it is very easy for concepts to play a useful role if we are liberal enough e.g. a concept can be useful because it is more familiar than an alternative, or more concise. This is a low bar to be satisfied, and one

that the Bayesian concept of confirmation satisfies—it is more concise to say 'E confirms H' than 'P(H|E&B) > P(H|B)'.[4] If it is good to use fewer concepts, then the Bayesian concept of confirmation is useful. Brössel and Huber say nothing against this use, so let's set it aside.

On the other hand, one could argue that the concept of confirmation is never needed because 'E confirms H' can always be replaced by: (D) P(E|H) > P(E|–H). An analysis of a concept shows that that concept is not fundamental. If it's best to use fewer fundamental concepts then an analysis is useful to the extent that it makes the concept dispensable—the concept itself can then be eliminated.

As it stands, this objection would apply to any conceptual analysis. Brössel and Huber focus only on the Bayesian concept of confirmation, so let's set aside this more general objection.

Still, there is a slightly different objection which applies only to disjunctive analyses and which seems to underlie Brössel and Huber's first objection. Suppose we have a disjunctive analysis of red:

S is red iff [S is scarlet or S is maroon]

A worry for this disjunctive analysis is that whether something is red is determined by whether it is scarlet or maroon, so redness cannot be used to determine whether an object is scarlet or maroon.

A similar worry seems to motivate Brössel and Huber. The Bayesian concept of confirmation is disjunctive, as there are many ways that E can confirm H, depending on the details of the probability function:

E confirms H iff [{P(E|H) = 1 > P(E|–H) = 0.5} or {P(E|H) = 1 > P(E|–H) = 0.6} or {P(E|H) = 0.9 > P(E|–H) = 0.5} or…]

The worry is that, '[s]ince the agent's degrees of belief are used to determine whether the evidence confirms the hypothesis, [1] confirmation cannot be used to determine the agent's degrees of belief, that is, how worthy of belief the hypothesis is' (Brössel and Huber 2015, 738).[5]

I will argue that we can make claims about confirmation without prior claims about how worthy of belief the hypothesis is (sections 4 and 5); and that even if we could not, the concept of confirmation could still be useful (section 6). A disjunctive analysis is more coarse-grained than the disjuncts, so two complementary purposes for coarse-grained concepts will emerge—they are useful when we are ignorant of the details and they are useful when omitting details improves an explanation.

A different objection one might have to a conceptual analysis is that it doesn't match the target concept sufficiently closely. This worry seems to motivate Brössel and Huber's objection to the second purpose they consider, that confirmation can be used to (2) measure the epistemic value of experimental outcomes. Brössel and Huber claim that this is incompatible with the desiderata that old evidence can confirm hypotheses. Specifically, they object to extant answers to the old evidence problem on the grounds that these answers are incompatible with new evidence confirming hypotheses. They argue that no theory of confirmation can account for both old evidence and new evidence. I will argue that familiar responses to the old evidence problem can also account for new evidence, and that even if they couldn't there would be still be a purpose for the Bayesian concept of confirmation (section 7).

Here's the plan. Section 3 explains Brössel and Huber's circularity objection. Section 4 argues that we can avoid circularity by defining credences and confirmation relations in different probability functions. Section 5 argues that even if we can't, coarse-grained confirmation relations are useful when we don't have detailed information about the beliefs of the agents we are describing. Section 6 argues that even if we do have detailed information about the beliefs of the agents we are describing, coarse-grained confirmation relations can be explanatorily useful. Section 7 argues that the Bayesian conception of confirmation can be useful for the purpose of (2) measuring the epistemic value of experimental outcomes.

3. The Circularity Problem

The bulk of Brössel and Huber's discussion—and mine—concerns (1). Let's start with Brössel and Huber's argument. They quote the following passage of Hempel:

> It is now clear that an analysis of confirmation is of fundamental importance also for the study of the central problem of what is customarily called epistemology; this problem may be characterized as the elaboration of 'standards of rational belief'. (Hempel 1945, 7)

And then they comment: 'We claim that no Bayesian conception of confirmation can be used for this purpose' (Brössel and Huber 2015, 740). The core of their argument is the following circularity problem:

> we must specify the agent's … degrees of belief before we can say whether the evidence confirms the hypothesis. Therefore we cannot use the information that the evidence confirms the hypothesis in order to specify the agent's actual degrees of belief [or what is worthy of belief]. (Brössel and Huber 2015, 740–741)

Specifically, E confirms H iff $P(E|H) > P(E|{-}H)$, where P expresses a rational belief function/credence function.[6] The beliefs must be settled first, and these determine the facts about confirmation. Thus confirmation is no help for making claims about how worthy of belief hypotheses are, as the order of explanation goes the other way. So Brössel and Huber object that as the Bayesian analysis defines confirmation in terms of degrees of belief, confirmation cannot be used to specify degrees of belief.

I will offer two reasons to think that *we need not specify the agent's degrees of* belief before we can say whether the evidence confirms the hypothesis. First, confirmation might not be defined in terms of degrees of beliefs (section 4). Second, we might know the (coarse-grained) confirmation relations without knowing the (fine-grained) degrees of belief (section 5). So my focus will be on the first sentence of the quotation above: 'we must specify the agent's…degrees of belief before we can say whether the evidence confirms the hypothesis' (Brössel and Huber 2015, 740; section 6 grants that the sentence is true.).

4. Alternative Probability Functions

I deny that 'we must specify the agent's … degrees of belief before we can say whether the evidence confirms the hypothesis' (Brössel and Huber 2015, 740). The reason is that confirmation and credence can be defined on different probability functions. For example, we can define confirmation in terms of physical probabilities i.e. chance.[7] Chances are physical features of the world that are separate from actual or ideal beliefs. Patrick Maher explains the difference:

> [S]uppose you have been told that a coin is either two-headed or two-tailed but you have no information about which it is. The coin is about to be tossed. What is the probability that it will land heads? There are two natural answers to this question:
>
> (i) 1/2.
>
> (ii) Either 0 or 1 but I do not know which.
>
> Answer (i) is natural if the question is taken to be about inductive probability [rational belief], while (ii) is the natural answer if the question is taken to be about physical probability [chance]. (Maher 2006, 185)

If confirmation can be understood in terms of chances, *we need not specify the agent's* degrees of belief before we can say whether the evidence confirms the hypothesis.

In fact explanations will often go the other way—we need to specify whether the evidence confirms the hypothesis before we can say what the agent's degrees of belief ought to be. For example, suppose an agent has been told that the coin is being tossed by a machine, M1, which is biased towards heads. Suppose the chance facts include:

> Ch(Heads | the coin is tossed by M1) > Ch(Heads | the coin is not tossed by M1)

This lets us define a confirmation relation in terms of chances: that the coin is tossed by M1 confirms Heads. Now we add that an agent knows that these are the chances, and matches their credences to the known chances (so we replace 'Ch' for chance with 'Cr' for credence):[8]

Cr(Heads | the coin is tossed by M1) > Cr(Heads | the coin is not tossed by M1)

Assuming the agent updates by conditionalization,[9] we need to specify at least the agent's prior credence in Heads, Cr(H) plus the known facts about confirmation above, to determine what her credence should be after learning M1. Thus the Bayesian conception of confirmation can be used for (1) making claims about how worthy of belief various hypotheses are.

This is a toy model of what often happens when people consider what to believe. When wondering, say, whether a Democrat will win the next election, we think about voting patterns, the economy, the appeal of the candidates etc. It is natural to say that we are trying to discern the objective chances, and using these to determine what we should believe.

Brössel and Huber might object that confirmation is defined only relative to a complete probability function, and chance does not have a complete probability function, e.g. there is no value for the unconditional chance that the coin is tossed by M1 (Hájek 2003, 296).[10]

One response is to modify the analysis of confirmation so it applies even when there is only a partial probability function. A second response is to point out that the appeal to chance is not essential to this strategy. All that's needed is that confirmation is defined in a probability function that differs from the probability function that represents the agent's beliefs. We could do this by defining confirmation in terms of inductive/evidential probability, which are different from the agent's actual subjective degrees of belief.[11]

Brössel and Huber consider this possibility, specifically, the proposals of Williamson (2000) and Hawthorne (2005) for such an inductive/evidential probability function, but reject them on the grounds that we have insufficient reason to believe they exist. But plenty of powerful reasons have been offered (Russell 1946, 646; White 2005; Maher 2006), not least that they solve the problem of induction, so their existence is at worst an open question. Still, perhaps Brössel and Huber are really interested in arguing that there is no purpose for a concept of confirmation defined in terms of degrees of belief. I'll use only this concept of confirmation for the rest of the paper.

5. Ignorance

Let's concede for the sake of argument that chance and inductive probabilities are not in good standing. Suppose we have only credence—subjective degrees of belief—to work with. Must we now specify the agent's degrees of belief before we can say whether the evidence confirms the hypothesis?

No. We might be ignorant of the full belief function of an agent, yet our partial knowledge of their belief function—as stated in terms of confirmation—helps specify what they should believe.[12]

For example, suppose we have the following information about an agent:

(A) The agent is probabilistically rational, updates by conditionalization, and at t_1 P(H) = 0.5. At t_2 they learn E.

Should P(H) at t_2 be more than 0.5? We don't have enough information to answer this question. Now add the following:

(B) E confirms H.

Should P(H) at t_2 be more than 0.5? We can now answer this question—the answer is 'yes'. Thus the Bayesian conception of confirmation can be used for (1) making claims about how worthy of belief various hypotheses are.

Our actual situation often has this shape. We more often know coarse-grained facts about confirmation according to an agent's credences than we know precise details of their credences. For example, let H = Einstein's theory of relativity, E = Eddington's photographs of the 1919 solar eclipse. It is plausible that we have the following information about an agent:

(A) The agent is probabilistically rational, updates by conditionalization, and in 1918 P(H) = 0.5. In 1919 they discover E.[13]

Should P(H) in 1919 be more than 0.5? We can answer 'yes' if we add:

(B) E confirms H.

So we can use the concept of confirmation to say something about the agent's degrees of belief, even if we don't know what her precise earlier degrees of belief were.

As well as making claims about individual agents' degrees of belief, we can also use the concept of confirmation to make claims about entire community's degrees of belief. We sometimes want to talk about the way a scientific community came to believe a new theory. If we are required to state the precise degrees of belief of every member of the community, we would never be able to do this. But putting things in terms of confirmation allows us to abstract away from the details of each individual, and make general claims about the community. If every member of the community had a credence function such that E confirms H, then we can conclude that E confirms H for the entire community.

Compare the analysis of red in terms of scarlet or maroon. We might be ignorant of whether an object is scarlet or maroon, but have the—very useful—information that it is red. And we might have a community of objects, some of which are scarlet and some maroon. The concept of red allows us to say useful things about the whole community.

6. Explanation

The last section discussed cases where we are ignorant of some details of the agent's credences. At this point Brössel and Huber might want to retreat to the more modest claim that the Bayesian concept of confirmation is useless when we have full information about the agent. That is, when we know exactly what the agent's credences are, there is nothing for the concept of confirmation to do. For example, suppose that as well as A, we are also told the specific likelihoods:

(C) $P(E|H) = 1$; $P(E|–H) = 0.25$

This would have allowed us to conclude that P(H) at t_2 is more than 0.5 without using the concept of confirmation.

I will argue that there is a role for the concept of confirmation even if we have full information. So I will grant the first sentence of the quote above: 'we must specify the agent's … degrees of belief before we can say whether the evidence confirms the hypothesis'. Nevertheless, there is still a role for the concept of confirmation; I will argue that:

> Even given full information about the agent's beliefs we can use the information that the evidence confirms the hypothesis in order to explain the agent's (rational[14]) degrees of belief.

This is based on the second sentence of the quotation ('Therefore we cannot use the information that the evidence confirms the hypothesis in order to specify the agent's degrees of belief'), but I've changed 'specify' to 'explain'.[15] I think that even if we have full information about the agent's beliefs, there is an explanatory role to be played by the concept of confirmation.

Repeating for convenience:

(B) E confirms H.
(C) $P(E|H) = 1$; $P(E|{-}H) = 0.25$.

The difference between B and C on which I want to focus is that B has fewer details than C. So the question is: assuming the agent learns (and conditionalises on) only E between t_1 and t_2, can an explanation of the value of P(H) at t_2 which uses B have any advantage over an explanation which uses C? That is, can Explanation 1 be better than Explanation 2?

> Explanation 1: $P_{t2}(H) > P_{t1}(H)$ because E confirms H (at t_1)
> Explanation 2: $P_{t2}(H) > P_{t1}(H)$ because $P_{t1}(E|H) = 1$ and $P_{t1}(E|{-}H) = 0.25$

Yes—there are many cases where omitting details improves an explanation. In Hilary Putnam's (1975) famous example, if we want to know why a square peg fails to fit in a

round hole, we are better off citing the logically weaker fact that the peg is square than the logically stronger description of every molecule of the peg. Alan Garfinkel (1981) argues that the best explanation of changes in rabbit populations do not make reference to the details of which rabbits were eaten by which foxes. Jerry A. Fodor (1987, 3–4) argues that if you want to explain his behaviour, you should work with his (high-level) desires and beliefs, rather than his (low-level) neurological states. And Jackson and Philip Pettit (1990) argue that a conductor's annoyance is better explained by the fact that someone is coughing than by the fact that Bob is coughing (assuming the conductor doesn't have a particular dislike of Bob). Such examples motivate various forms of functionalism and non-reductionism.

There are various theories about why removing details improves explanations. One theory is that explanations with fewer details are more robust across possibilities, and robustness is a virtue of explanations. In our case, the explanation in terms of confirmation is more robust, in that the explanation would survive even if the degrees of belief were slightly different. By contrast, an explanation that specifies the exact degrees of belief becomes false if the degrees of belief are at all different, so the purported explanation becomes false and fails as an explanation (White 2005; Jones 2018).[16] So an explanation with fewer details might be better for explaining more phenomena.

Another theory of why explanations with fewer details can be better is that explanation is contrastive, so the explanation has to be of the right level of generality to fit the explanandum (Schaffer 2005; Clarke 2016).[17] In our case, if we want to know why P(H) at t_2 is more, rather than less than at t_1, then we need to know whether E confirms or disconfirms H. We don't need to know whether P(E|–H) = 0.25 rather than 0.24. So if we are interested in what direction the agent's credences moved in, it is confirmation that matters, not the exact likelihoods. And indeed, we are often interested in explaining why scientists' confidence in a theory went up or down; we rarely care about the exact degrees of belief they had. So an explanation with fewer details might be better for showing the patterns of counterfactual dependence.

A third theory—one I like more—is that logically stronger explanations are better, and that the Explanation 1 is stronger because it omits details.[18] Explanation 2 tells us that in the specific situation where $P_{t1}(E|H) = 1$ and $P_{t1}(E|–H) = 0.25$, H becomes more probable. Explanation 1 tells us that in any[19] situation where E confirms

H, H becomes more probable. So Explanation 1 is logically stronger than Explanation 2 and this might be why Explanation 1 is better.

Any one of these accounts can be applied here, so we can remain neutral on which is correct. We ~~only~~ need only the assumption that omitting details can improve an explanation, and this is widely agreed.

To bring this together, Explanation 1 can be better than Explanation 2 in virtue of having fewer details. So even if we know all the facts about the agent's beliefs, we should not dispense with the concept of confirmation. The concept of confirmation allows us to state less detailed facts than those statable in terms of precise beliefs, and these less detailed facts can provide superior explanations of why agents have the beliefs they do.

Returning to the analysis of red in terms of scarlet and maroon, what is the purpose of the concept of red if we have full information about the shade of red? Suppose that bees are attracted to red. We might explain why a bee flew towards a flower by citing the flower's scarletness. But it is plausibly a better explanation to cite the flower's redness. Thus the disjunctive concept has a use in giving a superior explanation.

A referee has objected that Explanation 1 and Explanation 2 fail to be explanations at all. But I'm not sure on what grounds someone could deny it. There is a straightforward causal connection between $P_{t1}(E|H) > P_{t1}(E|{-}H)$, combined with conditionalising on E, and $P_{t2}(H) > P_{t1}(H)$. To put it in the terms of Hempel and Oppenheim (1948, 137), the 'antecedent conditions' are the t_1 probabilities and the 'general law' is conditionalization.

7. Confirmation Determines the Epistemic Value of Experimental Outcomes

Brössel and Huber consider and reject a different use for the Bayesian concept of confirmation—confirmation determines the epistemic value of experimental outcomes, and thus helps decide which experiments to carry out:

> One might argue that the epistemic value of an experimental outcome in a test of a hypothesis for an agent consists in the difference the experimental

> outcome would make to the agent's degree of belief in the hypothesis. In the literature, this is referred to as 'the potential further support' (Christensen 1999) or 'the additional evidence provided by' (Milne 2014, 255) an experimental outcome for a hypothesis for an agent. (Brössel and Huber 2015, 744; citations adjusted)

But Brössel and Huber reject this, arguing that most Bayesians want to give an analysis of actual support rather than potential further support, where actual support is provided by evidence the agent already has. And they argue that no analysis concerning actual evidence can be of any help regarding what potential future evidence we could look for. I will argue that Bayesians give an analysis of both actual and potential future support; the latter can help decide which experiments to perform.

First, why think that most Bayesians want to give an analysis of actual support? Because Bayesians hold that old evidence can still be confirming evidence (Glymour 1980). Old evidence is evidence the agent already has, so we can model it with: $P(E) = 1$. The old evidence problem is that old evidence cannot confirm anything, as $P(E) = 1$ trivially entails that $P(E|H) = P(E|–H) = 1$. In order to allow that old evidence can confirm hypotheses, Bayesians must complicate their conception of confirmation beyond 'E confirms H iff $P(E|H) > P(E|–H)$'. And once they've done so, they have an analysis of confirmation that includes actual support given by known evidence. This is all common ground between myself and Brössel and Huber.

But Brössel and Huber hold that an analysis of actual support cannot also be a measure of potential support. They give the following argument:

> To sum up: old evidence cannot provide incremental confirmation, and since potential further support is a form of incremental confirmation, old evidence cannot provide potential further support. Therefore all philosophers who take the problem of old evidence seriously cannot want to capture potential further support with their notions of confirmation. (Brössel and Huber 2015, 746)

But the second sentence does not follow from the first. Philosophers who take the problem of old evidence seriously can still want to capture potential further support with

their analysis of confirmation—they just need to have an analysis of confirmation that is not limited to old evidence.

Brössel and Huber claim that most Bayesians are ~~only~~ intending only to give an analysis of actual support. I'll first raise a simple problem for this suggestion. Suppose, for reductio, that Bayesians are ~~only~~ intending only to give an analysis of actual support. Now consider the following example. Consider the hypothesis, H, that a fair die will land showing a 6 on its next throw. Let E = The die will land showing an even number on its next throw. Does E confirm H according to Bayesian confirmation theory? Brössel and Huber must say no. For as E is about the next throw, E is merely potential evidence, and Brössel and Huber claim Bayesians only give an analysis of actual support. I take this to be an absurd result of the claim that Bayesians are only giving an analysis of actual support. Of course, Brössel and Huber might reply that that's all Bayesians can do, and so much the worse for Bayesianism.

I don't think that's all Bayesians can do, though. To demonstrate how we could have an analysis of confirmation that accounts for both actual and potential support, consider a flat-footed disjunctive analysis:

E confirms H iff either

i) E is unknown and $P(H|E) > P(H)$ or

ii) E is known and…

where we plug in our solution to the old evidence problem after the dots. Potential support is provided by (i) and actual support is provided by (ii). The problem of old evidence should be thought of as the problem of how to expand our analysis of confirmation to incorporate known evidence, not to replace our analysis of confirmation with one that ~~only~~ accounts only for known evidence.

I don't want to defend a disjunctive analysis; the point is just to show that an analysis of confirmation for both actual and potential support is possible. Indeed, we might hope for a confirmation theory that collapses the disjunction. For example, Brössel and Huber (2015, 741) quote ~~consider~~ Howson and Urbach:[20]

> the support of H by E is gauged according to the effect which one believes a knowledge of E would now have on one's degree of belief in H, on the

> (counterfactual) supposition that one does not yet know E. (Howson and Urbach 1993, 404–405, notation adapted; see also Howson and Urbach 2006, 297–301)~~. (Brössel and Huber 2015, 741)~~

As a first pass, this could be plugged into our schema as follows:

> E confirms H iff either
> i) E is unknown and P(H|E) > P(H) or
> ii) E is known and, on the supposition that one does not yet know E, P*(H|E) > P*(H) (where P* is the counterfactual credence function).

In the case where one really does not yet know E, supposing that one does not yet know E amounts to supposing what-is-known-to-be-actual. So when E is not yet known, the account collapses to the traditional analysis: E confirms H iff P(H|E) > P(H). To be clear, I am not defending Howson and Urbach's theory; my point is that it is open that there are non-disjunctive theories of confirmation that allow for actual and potential confirmation.

Brössel and Huber seem to assume that measures of confirmation must be functions of the probability function e.g. the r-measure, the d-measure, the l-measure, etc. (Fitelson 1999).~~[21]~~ And they also seem to assume that the same measure should account for both actual and potential support.[22] But we should reject both these assumptions. We are not limited to picking one of these measures in our efforts to solve the old evidence problem. We can take a more imaginative approach, such as appealing to counterfactuals, or to logical learning, or to something else.

In support of their view that Bayesians are ~~only~~ interested only in actual support, Brössel and Huber (2015, 745) quote Christensen, who wants

> to capture the support an agent's confidence in H already receives from E (in contrast to the potential further support that might be gotten from raising Pr(E)). (Christensen 1999, 449; notation adapted)

It is true that Christensen wants to do this, but that isn't all he wants to do. That quote is taken from a discussion that is trying to motivate a particular solution to the old

evidence problem. And Christensen motivates the proposed solution by explaining a little earlier that what we want in a solution to the old evidence problem is 'a measure of confirmation that goes beyond measuring potential further support' (Christensen 1999, 449; emphasis added). Christensen doesn't want a measure of confirmation that ignores potential support: he wants a measure that includes potential support and more besides. I think most Bayesians are the same.

Brössel and Huber might be motivated by the thought that none of the solutions to the old evidence problem succeed, and I have no interest in defending any of them. Let's grant that the old evidence problem cannot be solved. And let's grant that it follows that there is no precise concept of confirmation that matches the ordinary language concept of confirmation. A fortiori, the Bayesian concept of confirmation would fail to match the ordinary language concept of confirmation. But this is very different from Brössel and Huber's claim that the Bayesian concept of confirmation has no use. In fact, the opposite conclusion would be established. We would be left with the Bayesian concept of 'confirmation' as a very useful measure of potential support—and useless as a measure of actual support. It would just be a concept that poorly matches the ordinary language concept of confirmation.

Still, this would be no great cost, as the Bayesian conception fails to match the ordinary language concept very well. The ordinary language concept requires that E confirms H only if P(H|E) is high.[23] So the Bayesian concept of confirmation already departs from the ordinary language of confirmation in a significant way. The point here is that this does little to undermine its usefulness.

8. Conclusion

Brössel and Huber raise the question of what the Bayesian conception of confirmation is for. I have argued that it can be used for both the purposes considered by Brössel and Huber. First, it can explain what beliefs an agent ought to have, especially when we don't know the full facts, such as details of the agent's beliefs. Second, it can inform how valuable particular experiments are and, therefore, which should be performed.

Much of the discussion was specific to confirmation, but there are two more general morals. First, coarse-grained concepts can be more useful than fine-grained

concepts when i) we don't have full information or ii) omitting details improves an explanation. Second, concepts can be useful even when they differ from the ordinary language concept they are based on. We might end up closer to defining a new concept rather than analysing or explicating an existing one, but this does little to undermine the usefulness of the resulting concept.

Acknowledgements

I'm grateful to three referees for this journal for comments on earlier drafts of this paper.

References

Bradley, D. Forthcoming. "Should Explanations Omit the Details?" British Journal for the Philosophy of Science.

Brössel, P., and F. Huber. 2015. "Bayesian Confirmation: A Means with No End." British Journal for the Philosophy of Science 66: 737–749.

Carnap, R. 1962. Logical Foundations of Probability. 2nd ed. Chicago, IL: University of Chicago Press.

Carnap, R.~~, & Schilpp, P. A. (~~ 1963~~)~~.- "Replies and Systematic Expositions." In The Philosophy of Rudolf Carnap~~.~~, edited by P. A. Schilpp, 859–1013. LaSalle, IL: Open Court~~Cambridge University Press~~.

Christensen, D. 1999. "Measuring Confirmation." Journal of Philosophy 96: 437–461.

Clarke, C. 2016. "The Explanatory Virtue of Abstracting Away from Idiosyncratic and Messy Detail." Philosophical Studies 173: 1429–1449.

Fitelson, B. 1999. "The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity." Philosophy of Science 66 (Proceedings): S362–S378.

Fitelson, B., and A. Hájek. 2017. "Declarations of Independence." Synthese 194: 3979–3995.

Fodor, Jerry A. 1987. Psychosemantics: The Problem of Meaning in the Philosophy of Mind. Cambridge, MA: MIT Press.

Garfinkel, A. 1981. Forms of Explanation: Rethinking the Questions in Social Theory. New Haven, CT: Yale University Press.

Glymour, C. 1980. Theory and Evidence. Princeton, NJ: Princeton University Press.

Hájek, A. 2003. "What Conditional Probability Could Not Be." Synthese 137: 273–323.

Harman, G. 1986. Change in View: Principles of Reasoning. Cambridge, MA: MIT Press.

Hawthorne, J. 2005. "Degree-of-belief and Degree-of-support: Why Bayesians Need Both Notions." Mind 114: 277–320.

Hempel, C. G. 1945. "Studies in the Logic of Confirmation (I)." Mind 54: 1–26.

Hempel, C. G., and P. Oppenheim. 1948. "Studies in the Logic of Explanation." Philosophy of Science 15: 135–175.

Howson, C., and P. Urbach. 2006. Scientific Reasoning: The Bayesian Approach. 3rd ed. LaSalle, IL: Open Court.

Jackson, F., and P. Pettit. 1990. "Program Explanation: A General Perspective." Analysis 50: 107–117.

Jones, N. 2018. "Inference to the More Robust Explanation." British Journal for the Philosophy of Science 69: 75–102.

Lewis, D. 1980. "A Subjectivist's Guide to Objective Chance." In Ifs: Conditionals, Belief, Decision, Chance, and Time, edited by W. L. Harper, R. Stalnaker, and G. Pearce, (pp. 267–297). Dordrecht: SpringerD. Reidel.

Maher, P. 2006. "The Concept of Inductive Probability." Erkenntnis 65: 185–206.

Maher, P. 2007. "Explication Defended." Studia Logica 86: 331–341.

Milne, P. 2014. "Information, Confirmation, and Conditionals." Journal of Applied Logic 12, Issue 3: 252–262.

Putnam, H. 1975. "Philosophy and Our Mental Life." In H. Putnam, Mind, Language, and Reality: Philosophical Papers, Volume 2 vol. 2, 291–303. Cambridge: Cambridge University Press.

Russell, B. 1946. A History of Western Philosophy. London: George Allen and Unwin.

Salmon, W. C. 1975. "Confirmation and Relevance.'—" In Induction, Probability, and Confirmation, edited by G. Maxwell and R. M. Anderson, 3–36. Minneapolis: University of Minnesota PressMinnesota Studies in the Philosophy of Science, 6, 3-36.

Schaffer, J. 2005. "Contrastive Causation." Philosophical Review 114: 327–358.

Schroeder, M. 2011. "Ought, Agents, and Actions." Philosophical Review 120: 1–41.

Schwitzgebel, E. 2011. Perplexities of Consciousness. Cambridge, MA: MIT Press.

White, R. 2005. "Explanation as a Guide to Induction." *Philosophers' Imprint* 5(2): 1–29.

Williamson, T. 2000. Knowledge and Its Limits. Oxford: Oxford University Press.

Notes

[1] Salmon (1975). This is equivalent to P(E|H) > P(E|–H) assuming that 0 < P(H) < 1, so I will treat them as equivalent. See Fitelson and Hájek (2017).

[2] We have more plausibly an explication than an analysis. Indeed this is the example that motivated Carnap to make the distinction; see Carnap ~~and Schilpp~~ (1963, 933–940) and Maher (2007). I'll address the question of whether we are analysing, improving or replacing the ordinary language concept at the end of section 7. Of course, there are alternative explications of confirmation beyond the Bayesian explication. One question for Brössel and Huber is whether they think their analysis undermines these other concepts of confirmation, or whether the Bayesian concept has a special problem.

[3] Thanks to a referee for comments that led to me setting the issues up in this way, and for various other improvements.

[4] I use quotation marks and inverted commas only to help with parsing; the distinction between sentences and propositions will be play no role.

[5] This passage shows that it is not entirely clear whether Brössel and Huber are working with actual or ideal beliefs. They say that they are focusing on 'Bayesian confirmation theory qua normative theory' (Brössel and Huber 2015, 738n1). And in order to engage with Hempel, who talks about rational belief, it must be what agents ought to believe that is at issue. But they reject the evidential/inductive probability functions that rationality seems to require (Brössel and Huber 2015, 743). So where does the normativity come from? I think they must have in mind a theory according to which ~~conditionalization~~ conditionalisation is the only normative constraint (plus probabilism). And I will assume for simplicity that the agents we are dealing with are sufficiently ideal for the actual/ideal distinction to collapse.

[6] It is sometimes more natural to talk about credences, but for the most part I will follow Brössel and Huber and talk about beliefs.

[7] Brössel and Huber consider using different probability functions for credence and confirmation, but don't consider objective chance.

[8] Something like the Principal Principle (Lewis 1980) is needed, though a much weaker version will do the job regarding conditional probabilities.

[9] Conditionalization says that if an agent learns exactly E between t_1 and t_2, then $P_{t1}(H) = P_{t2}(H|E)$.

[10] ~~Hájek (2003), 296.~~

[11] They are different even if the agent is ideal. Deductive logic constrains what agents should believe, but substantive bridging principles are needed to connect deductive logic with belief. Similarly, inductive logic constrains what agents should believe, and substantive bridging principles are also needed here. See Harman (1986).

[12] Indeed we might be ignorant of our own beliefs (Williamson 2000; Schwitzgebel 2011), and perhaps have better access to coarse-grained confirmation relations.

[13] I'm using E for both the photographs and the proposition that would be learnt on seeing them.

[14] Again, it is rational degrees of belief that are relevant for Brössel and Huber's claim that the Bayesian conception of confirmation is useless for making claims about how 'worthy of belief various hypotheses are'.

[15] The original sentence is: 'Therefore we cannot use the information that the evidence confirms the hypothesis in order to specify the agent's degrees of belief'.

[16] ~~White (2005); Jones (2018).~~

[17] ~~Schaffer (2005); Clarke (2016).~~

[18] This might look paradoxical—but omitting details from the antecedent of a conditional makes it stronger. See Bradley (Forthc.).

[19] We might need to add a ceteris paribus clause here.

[20] And here is this disjunctive schema with Garber's (1983) logical learning solution to the old evidence problem plugged in: E confirms H iff either i) E is unknown and P(H|E) > P(H) or ii) E is known and P(H|H entails E) > P(H).

[21] ~~Fitelson (1999)~~

[22] Brössel and Huber (2015~~, 746~~) seem to make these assumptions in the middle of p. 746.

[23] Incremental and absolute confirmation were conflated by Hempel (1945) and clarified by Carnap (1962, 477–478).