

What Is the Point of Confirmation?

Franz Huber^{†‡}

Philosophically, one of the most important questions in the enterprise termed confirmation theory is this: Why should one stick to well confirmed theories rather than to any other theories? This paper discusses the answers to this question one gets from absolute and incremental Bayesian confirmation theory. According to absolute confirmation, one should accept “absolutely well confirmed” theories, because absolute confirmation takes one to true theories. An examination of two popular measures of incremental confirmation suggests the view that one should stick to incrementally well confirmed theories, because incremental confirmation takes one to (the most) informative (among all) true theories. However, incremental confirmation does not further this goal in general. I close by presenting a necessary and sufficient condition for revealing the confirmational structure in almost every world when presented separating data.

1. Introduction. Philosophically, one of the most important questions in the enterprise traditionally termed confirmation theory is this: Why should one stick to well confirmed theories rather than to any other theories? In other and more mundane words: What is the point of confirmation? In what follows I will examine whether and how absolute and incremental Bayesian confirmation theory answer this question.

According to absolute Bayesian confirmation theory, an agent’s *degree of absolute confirmation* of some hypothesis or theory H by a piece of evidence E relative to a body of background information B equals the probability of H given E and B , $\Pr(H|E \wedge B)$, where $\Pr : \mathcal{L} \rightarrow \mathfrak{R}$ is the agent’s actual degree of belief function on some language \mathcal{L} (see Section 2). According to incremental Bayesian confirmation theory, an agent’s

To contact the author, please write to: Philosophy, Probability, and Modeling Research Group, Center for Junior Research Fellows, University of Konstanz, P.O. Box M 682, D-78457 Konstanz, Germany; e-mail: franz.huber@uni-konstanz.de.

^{†‡}I have benefited from discussions with Luc Bovens, Branden Fitelson, Jim Joyce, Sherri Roush, and Gerhard Schurz. My research was supported by the Alexander von Humboldt Foundation, the Federal Ministry of Education and Research, and the Program for the Investment in the Future (ZIP) of the German Government through a Sofja Kovalevskaja Award.

Philosophy of Science, 72 (December 2005) pp. 1146–1159. 0031-8248/2005/7205-0041\$10.00
Copyright 2005 by the Philosophy of Science Association. All rights reserved.

degree of incremental confirmation of H by E relative to B is measured by a relevance measure r_{Pr} based on the agent's actual degree of belief function Pr ; i.e., a possibly partial function $r_{Pr} : \mathcal{L} \times \mathcal{L} \times \mathcal{L} \rightarrow \mathfrak{R}$ such that for all $H, E, B \in \mathcal{L}$ with $Pr(E \wedge B) > 0$:

$$r_{Pr}(H, E, B) \cong 0 \Leftrightarrow Pr(H|E \wedge B) \cong Pr(H|B).$$

2. The Point of Absolute Confirmation. The traditional answer to our question is something like this: Science aims at true theories, and one should accept well confirmed theories, because confirmation takes one to true theories. Indeed, if arriving at true theories is our (only) goal, then there is a point to absolute confirmation. In the long run, absolute confirmation almost surely takes one to true theories. This is the content of the following theorem (Gaifman and Snir 1982, 507):

Theorem 1 (Gaifman and Snir). Let $S = \{A_i \in \mathcal{L} : i = 0, 1, \dots\}$ separate Mod_L , let A_i^ω be A_i if $\omega \models A_i$ and $\neg A_i$ otherwise, and let $[B](\omega)$ be 1 if $\omega \models B$ and 0 otherwise. Then for every $B \in \mathcal{L}$,

$$Pr\left(B \mid \bigwedge_{0 \leq i < n} A_i^\omega\right) \rightarrow [B](\omega) \text{ almost everywhere as } n \rightarrow \infty.$$

Here is the relevant technical background. L is obtained from a first-order language for arithmetic, L_0 , by adding finitely many "empirical" predicates and function symbols (whose interpretation is not fixed). L_0 contains all numerals '1', ... as individual constants; countably many individual variables ' x_1 ', ... taking values in the set of natural numbers N ; the common symbols '+', ' \cdot ', and '=' for addition, multiplication, and identity, respectively; and the standard quantifiers and connectives. In addition, there may be finitely many predicates and function symbols denoting certain fixed relations over N . The set of well formed formulas of L is denoted by ' \mathcal{L} ' and is also called a language.

A model for L consists of an interpretation φ of the empirical symbols which assigns every k -ary predicate ' P ' a subset $\varphi(P) \subseteq N^k$, and every k -ary function symbol ' f ' a function $\varphi(f)$ from N^k to N . The interpretation of the symbols in L_0 is the standard one and is kept the same in all models. Mod_L is the set of all models for L . ' $\omega \models A$ ' says that formula A is true in model $\omega \in \text{Mod}_L$. $A[x_1, \dots, x_k]$ is valid, $\models A[x_1, \dots, x_k]$, iff $\omega \models A[n_1/x_1, \dots, n_k/x_k]$ for all $\omega \in \text{Mod}_L$ and all $n_1, \dots, n_k \in N$. Here, ' $A[n_1/x_1, \dots, n_k/x_k]$ ' results from ' $A[x_1, \dots, x_k]$ ' by uniformly substituting ' n_i ' for ' x_i ' in ' A ', $1 \leq i \leq k$. ' $A[x_1, \dots, x_k]$ ' indicates that ' x_1, \dots, x_k ' are the only variables occurring free in ' A '.

A function $Pr : \mathcal{L} \rightarrow \mathfrak{R}_{\geq 0}$ is a *probability on \mathcal{L}* iff for all $A, B \in \mathcal{L}$:

1. $\models A \leftrightarrow B \Rightarrow \Pr(A) = \Pr(B)$
2. $\models A \Rightarrow \Pr(A) = 1$
3. $\models \neg(A \wedge B) \Rightarrow \Pr(A \vee B) = \Pr(A) + \Pr(B)$
4. $\Pr(\exists x A[x]) = \sup \{ \Pr(A[n_1/x] \vee \dots \vee A[n_k/x]) : n_1, \dots, n_k, k \in \mathbb{N} \}$

The conditional probability of A given B , $\Pr(A|B)$, is defined as

5. $\Pr(A|B) = \Pr(A \wedge B) / \Pr(B)$

provided $\Pr(B) > 0$. \Pr is *regular* iff the converse of 2. holds as well,

6. $\Pr(A) = 1 \Rightarrow \models A$.

A set of sentences $S \subseteq \mathcal{L}$ *separates* a set of models $X \subseteq \text{Mod}_L$ iff for any two distinct $\omega_1, \omega_2 \in X$ there exists $A \in S$ such that $\omega_1 \models A$ and $\omega_2 \not\models A$. The set of all atomic empirical sentences separates Mod_L (Gaifman and Snir 1982, 507).¹

However, absolute confirmation has long been abandoned in favor of incremental confirmation. Is there another goal for incremental confirmation that is different from arriving at true theories? If so, what is this goal?

3. What Is the Point of Incremental Confirmation? Two popular measures of incremental confirmation are the distance measure d (Earman 1992) and the Joyce-Christensen measure s (Joyce 1999; Christensen 1999):

$$d_{\Pr}(H, E, B) = \Pr(H|E \wedge B) - \Pr(H|B),$$

$$s_{\Pr}(H, E, B) = \Pr(H|E \wedge B) - \Pr(H|\neg E \wedge B).$$

What do these measures measure? The measure d increases with

- the *plausibility* of H given E and B , $p = \Pr(H|E \wedge B)$, and
- the evidence neglecting or data independent semantic *informativeness* of H relative to B , $i_0 = \Pr(\neg H|B)$.

Similarly, s increases with

- the plausibility of H given E and B , $p = \Pr(H|E \wedge B)$, and

1. The Gaifman and Snir framework is not rich enough for proper theory assessment. The reason is that the “theories” whose truth values one converges to by conditioning on appropriate data sentences are formulated within the same “empirical” vocabulary as are the data sentences. So there is no room for theoretical terms in the sense that the probability of a theory whose formulation contains theoretical terms not occurring in any data sentence does not necessarily converge to its truth value when one keeps conditionalizing on these data sentences. As an aside, note that this problem disappears if the realist goal of truth is replaced by the empiricist goal of empirical adequacy.

- the evidence based or data dependent semantic informativeness of H relative to E and B , i.e., the amount to which H informs about E relative to B , $i_1 = \Pr(\neg H|\neg E \wedge B)$.

This is clearly seen by rewriting d and s as follows:

$$d_{\text{Pr}}(H, E, B) = \Pr(H|E \wedge B) + \Pr(\neg H|B) - 1,$$

$$s_{\text{Pr}}(H, E, B) = \Pr(H|E \wedge B) + \Pr(\neg H|\neg E \wedge B) - 1.$$

The quantities p and i_0 as well as p and i_1 are *conflicting* in the sense that p decreases, whereas i_0 and i_1 increase with the logical strength of the hypothesis to be assessed. So d and s weigh between two conflicting aspects, viz. the plausibility and the informativeness of the hypothesis to be assessed.

In Section 4 I will argue in more detail that i_0 and i_1 measure two different, but equally sensible kinds of informativeness. Section 5 provides another argument that: (i) d and s do nothing but weigh between the two conflicting goals of plausibility and informativeness; (ii) that they are exactly alike in the way they weigh between these two aspects; and (iii) that they differ from each other just in the respect that d is based on data independent informativeness whereas s is based on informativeness about the data. All this suggests the following answer to our question: Science aims at informative true theories, and one should stick to incrementally well confirmed theories, because incremental confirmation takes one to (the most) informative (among all) true theories. However, as shown in Section 6, incremental confirmation does not further this goal in general. I close by giving a necessary and sufficient condition for revealing the confirmational structure in almost every world when presented separating data.

4. Measuring Semantic Information. In a subjective Bayesian framework it is clear that $p = \Pr(H|E \wedge B)$ measures the *plausibility* of H in view of E and B . It is still rather obvious that $i_0 = \Pr(\neg H|B)$ measures the data independent informativeness of H relative to B . i_0 was already considered by Carnap and Bar-Hillel (1952), Bar-Hillel and Carnap (1953), Hempel (1960, 1962), and Hintikka and Pietarinen (1966) (for the notion of semantic information cf. Bar-Hillel 1952, 1955). The second measure that was discussed in this connection is

$$i_2 = -\log_2 \Pr(H|B) = \log_2 \frac{1}{\Pr(H|B)}.$$

i_2 is ordinally equivalent to i_0 . For future reference it is convenient to




Figure 1.

define the analogous

$$p_2 = \log_2 \Pr(H|E \wedge B),$$

which is ordinally equivalent to $p_0 = p$.

It is less obvious that $i_1 = \Pr(\neg H|\neg E \wedge B)$ measures how much H informs about the data E relative to background B (cf., however, Hilpinen 1970). Following the above mentioned literature, one would expect something like:²

$$i_3 = \Pr(\neg H|E \wedge B),$$

$$i_4 = \Pr(E) \cdot \Pr(\neg H|E \wedge B),$$

$$i_5 = \log_2 \frac{1}{\Pr(H|E \wedge B)} = -\log_2 \Pr(H|E \wedge B).$$

As is often the case, a picture is worth a 1,000 words—see Figure 1.

The background information B determines the set of possibilities and is nothing but a restriction on the set of possible worlds over which inquiry has to succeed. H is the hypothesis whose informativeness about the data E is to be assessed (relative to B). Suppose you are asked to strengthen H by deleting possibilities verifying it, that is, by shrinking the area representing H . Would you not delete possibilities outside E ? After all, given E , those are exactly the possibilities known not to be the actual one, whereas those possibilities inside E are still alive options. Indeed, i_1 increases when H shrinks to H' as depicted in Figure 2, because it measures how much of $\neg E$ is occupied by $\neg H$. As a consequence, the information

2. In Levi (1967), i_3 is proposed as, roughly, a measure for the relief from agnosticism afforded by accepting H as strongest relative to total evidence $E \wedge B$. For i_4 and i_5 the reader is referred to Hintikka and Pietarinen (1966).




Figure 2.

H provides about E is maximal if H logically implies E (in this case H is completely within E , and so $\neg H$ covers all of $\neg E$). So according to i_1 , two hypotheses both logically implying all of the data—say, a complete theory about the world, and a theory-like collection of the data—carry the same maximal amount of information about E . In a sense, this is odd, because one would like the complete theory to come out as more informative than the theory-like collection of the data. This is what i_0 yields. For i_0 it does not matter which possibilities one deletes in strengthening H (provided all possibilities have equal weight on the probability measure Pr). i_0 neglects whether they are inside or outside E . The other candidates for measuring semantic information do rather poorly on this count: they require the deletion of the possibilities inside E . (Another reason why i_3 , i_4 , and i_5 seem to be inappropriate in the present context is presented in the next section.)

The background information B plays a role different from that of the evidence E for i_0 and i_1 , but not for i_3 , i_4 , or i_5 . Clearly, there is a difference between data on the one hand and background assumptions on the other; and this difference should show up somewhere. Apart from the above mentioned point that B determines the set of possibilities over which inquiry has to succeed, whereas E is gathered in order to indicate which of these possibilities is the actual one, there is the following difference: Hypotheses are supposed to inform about the world, and hence about the data, but they are usually not supposed to inform about the background assumptions. (If one holds there should be no difference between E and B as far as measuring information is concerned, then one can nevertheless adopt the above measures by substituting $E' = E \wedge B$ and $B' = \top$ for E and B , respectively.)

In order to avoid that one has to take sides between i_0 and i_1 let us call a possibly partial function $i = f_{i_0, i_1} : \mathcal{L} \times \mathcal{L} \times \mathcal{L} \rightarrow [0, 1]$ a *strength*

indicator (based on i_0 and i_1) iff f is non-decreasing in both and increasing in at least one of its arguments i_0 and i_1 , $f_{i_0, i_1} = 1$ for $i_0 = i_1 = 1$, and $f_{i_0, i_1} = 0$ for $i_0 = i_1 = 0$.

5. Expected Informativeness as One Way of Weighing. Having tried to make plausible that i_0 and i_1 measure informativeness *per se* and informativeness about the data, respectively, let us now turn back to the distance measure d and the Joyce-Christensen measure s . The two conflicting goals of informativeness and plausibility are equally important for d and s —and they are all that matters for them. Hence, other things being equal—these other things being the probabilities (plausibility values) of the hypotheses given the data E and the background information B —the overall d - or s -value of hypothesis H relative to E and B is the greater, the higher the informativeness of H (in the respective sense).

Clearly, if one knows the truth values of the theories one is assessing, then the plausibility of a theory's being true is of no interest anymore. In this case all that matters is how informative the theories are. Yet in general we do not know these truth values. Hence we consider how plausible it is that they are true in the world we are in, and how informative they are (about this world). Then we form their overall value by combining these two parameters in some suitable way. One such way immediately suggests itself: assign H as its overall value its *expected informativeness*.

$$\begin{aligned} E(i_0) &= \Pr(\neg H|B) \cdot \Pr(H|E \wedge B) \\ &\quad - \Pr(\neg\neg H|B) \cdot \Pr(\neg H|E \wedge B), \\ E(i_1) &= \Pr(\neg H|\neg E \wedge B) \cdot \Pr(H|E \wedge B) \\ &\quad - \Pr(\neg\neg H|\neg E \wedge B) \cdot \Pr(\neg H|E \wedge B). \end{aligned}$$

A little bit of reformulation shows that

$$E(i_0) = d_{\Pr}(H, E, B) \quad \text{and} \quad E(i_1) = s_{\Pr}(H, E, B).$$

So once again, d and s are exactly alike in the way they combine or weigh between informativeness and plausibility—which is to form the expected informativeness (cf. Hintikka and Pietarinen 1966 and Levi 1961, 1963, but also Hempel 1960). Their sole difference lies in the way they measure informativeness. In this sense, part of the discussion about the right measure of incremental confirmation is a discussion about the right measure of semantic information.

The measures i_3 , i_4 , and i_5 do again poorly:

$$E(i_3) = E(i_4) = 0,$$

$$E(i_5) \cong 0 \Leftrightarrow \Pr(H|E \wedge B) \cong \Pr(\neg H|E \wedge B).$$

Hence only i_5 gives a non-trivial answer, viz. to maximize probability. But then we can simply stick to probabilities and need not employ i_5 .

6. Revealing the Confirmational Structure. The preceding suggests the following answer to the question what goal incremental confirmation is supposed to further: Science aims at informative truth, and one should stick to incrementally well confirmed theories, because incremental confirmation takes one to (the most) informative (among all) true theories. The question is, of course, whether and in what sense this holds true.

When is one theory at least as informative as another? Well, if the first theory logically implies the second one, then the first theory is at least as informative as the second one. When else? In general, there is no further condition that applies equally to all probability measures \Pr . Just as the only \Pr -independent condition for H_1 to be at least as probable as H_2 is that H_2 logically implies H_1 , so the above is the only \Pr -independent condition for H_1 to be at least as informative as H_2 .

Hence, given a possible world $\omega \in \text{Mod}(B)$, H_1 is to be preferred over H_2 in ω if H_1 is true in ω , but H_2 is false in ω ; or if H_1 and H_2 have the same truth value in ω , and H_1 logically implies H_2 but H_2 does not logically imply H_1 . If H is logically true, then H is preferred in ω over any H_2 which is false in ω . On the other hand, any contingent H_1 that is true in ω is preferred over H , because these H_1 's are not only true in ω ; they are also more informative than H . Similarly, if H is logically false, then H is worse in ω than any theory that is true in ω , but better than any theory that is false in ω (because they are all less informative than H).

In this way each ω induces a partial order among the set of all (equivalence classes of axiomatizations of) theories: On the positive side one has all theories that are contingently true in ω , and on the negative side there are all theories that are contingently false in ω . In between there are the logically determined theories. Among the true theories on the positive side, the most informative, i.e., the complete theory about ω , is on top, followed by all true hypotheses it logically implies, partially ordered according to the logical consequence relation. This order goes all the way down to the least informative among all true theories, the tautology, which is placed at the bottom of the positive side. On that same level is the most informative among all false theories, the contradiction, followed by all contingently false theories, again partially ordered ac-

ording to the logical consequence relation. Let us call this partial order *the confirmational structure of ω* .

For a given ω , we would like a function f to *stabilize* to the correct answer in the sense that f gets the confirmational structure of ω right after finitely many steps (data sentences from ω), and continues to do so forever without necessarily halting (or giving any other sign that it has arrived at the true answer)—cf. Kelly (1996). In general, stabilisation to the correct answer is a stronger requirement than convergence to the correct answer. However, the Gaifman and Snir convergence theorem actually gives rise to a measure 1 stabilisation result (assign 1 to H if its probability exceeds .5, and 0 otherwise).

Let e_0, \dots, e_n, \dots be a sequence of sentences all of which are true in $\omega \in \text{Mod}(B)$. A possibly partial function $f: \mathcal{L} \times \mathcal{L} \times \mathcal{L} \rightarrow \mathfrak{R}$ *reveals the confirmational structure of ω when presented $(e_i)_{i \in \mathbb{N}}$* iff for any contingent $H_1, H_2 \in \mathcal{L}$, and any $H \in \mathcal{L}$:

1. $\omega \models H_1, \omega \not\models H_2 \Rightarrow \exists n \forall m \geq n: f(H_1, E_m, B) > 0 > f(H_2, E_m, B)$
2. $\omega \models H_1, \omega \models H_2, H_1 \models H_2 \neq H_1 \Rightarrow \exists n \forall m \geq n: f(H_1, E_m, B) > f(H_2, E_m, B) > 0$
3. $\omega \not\models H_1, \omega \not\models H_2, H_1 \models H_2 \neq H_1 \Rightarrow \exists n \forall m \geq n: 0 > f(H_1, E_m, B) > f(H_2, E_m, B)$
4. $\models H$ or $\models \neg H \Rightarrow \forall m: f(H, E_m, B) = 0$,

where $E_m = \bigwedge_{0 \leq i < m} e_i$. An immediate consequence of the Gaifman and Snir convergence theorem is

Observation 1. For any regular Pr on \mathcal{L} and any $\{e_i \in \mathcal{L} : i \in \mathbb{N}\}$ separating Mod_L there is $X \subseteq \text{Mod}_L$ with $\text{Pr}^*(X) = 1$ such that for all $\omega \in X$ (and hence for all $\omega \in X \cap \text{Mod}(B)$, for any $B \in \mathcal{L}$): d_{Pr} , s_{Pr} , and c_{Pr} it reveal the confirmational structure of ω when presented $(e_i)_{i \in \mathbb{N}}$.

Pr^* is the unique probability measure on the smallest σ -field \mathcal{A} containing the field $\{\text{Mod}(A) : A \in \mathcal{L}\}$ such that $\text{Pr}(A) = \text{Pr}^*(\text{Mod}(A))$ for all $A \in \mathcal{L}$. c is the Carnap measure (Carnap 1962),

$$\begin{aligned} c_{\text{Pr}}(H, E, B) &= \text{Pr}(H \wedge E \wedge B) \cdot \text{Pr}(B) - \text{Pr}(H \wedge B) \cdot \text{Pr}(E \wedge B) \\ &= (p + i_0 - 1) \cdot \text{Pr}(B) \cdot \text{Pr}(E \wedge B). \end{aligned}$$

However, observation 1 does not extend to all relevance measures. The log-ratio measure r (Milne 1996) and the log-likelihood ratio measure l (Fitelson 1999, 2001) do not reveal the confirmational structure of almost

every $\omega \in \text{Mod}_L$ when presented separating data:

$$r_{\text{Pr}}(H, E, B) = \log \left[\frac{\text{Pr}(H|E \wedge B)}{\text{Pr}(H|B)} \right],$$

$$l_{\text{Pr}}(H, E, B) = \log \left[\frac{\text{Pr}(E|H \wedge B)}{\text{Pr}(E|\neg H \wedge B)} \right]$$

$$= \log \left[\frac{\text{Pr}(H|E \wedge B) \cdot \text{Pr}(\neg H|B)}{\text{Pr}(\neg H|E \wedge B) \cdot \text{Pr}(H|B)} \right].$$

Like all relevance measures, r and l separate contingently true from contingently false theories. More precisely, for any regular Pr on \mathcal{L} , any $\{e_i \in \mathcal{L} : i \in N\}$ separating Mod_L , any $B \in \mathcal{L}$, any $\omega \in X \cap \text{Mod}(B)$ (for some $X \subseteq \text{Mod}_L$ with $\text{Pr}^*(X) = 1$), and any two contingent $H_1, H_2 \in \mathcal{L}$ such that $\omega \models H_1$ and $\omega \not\models H_2$ there exists n such that for all $m \geq n$,

$$r_{\text{Pr}}(H_1, E_m^\omega, B) > 0 > r_{\text{Pr}}(H_2, E_m^\omega, B), \quad r = r, l.$$

Furthermore, r and l also weigh between plausibility and informativeness:

$$r_{\text{Pr}}(H, E, B) = \log \text{Pr}(H|E \wedge B) - \log \text{Pr}(H|B)$$

$$= p_2 + i_2,$$

$$l_{\text{Pr}}(H, E, B) = \log \text{Pr}(H|E \wedge B) - \log \text{Pr}(H|B)$$

$$- (\log \text{Pr}(\neg H|E \wedge B) - \log \text{Pr}(\neg H|B))$$

$$= p_2(H) + i_2(H) - (p_2(\neg H) + i_2(\neg H)).$$

However, although r does distinguish between informative and uninformative true theories (in the sense of revealing part 2 of the confirmational structure of almost every world), it does not distinguish between informative and uninformative false theories. l performs even worse on this count, because it neither distinguishes between informative and uninformative true theories nor between informative and uninformative false theories. The reason is fairly obvious: If $p = 0$, then $p_2 = \log p = -\infty$, whence $p_2 + i_2 = -\infty$ for any finite value of i_2 . This means in particular that informativeness does not matter anymore once a theory is falsified by the data. Similarly in case of r .

Which conditions are sufficient for a function to reveal the confirmational structure of almost every world when presented separating data? Let $f = f(i, p)$ be a function of, among others, $p = \text{Pr}(H|E \wedge B)$ and some strength indicator $i = f_{i_0, i_1}$ based on $i_0 = \text{Pr}(\neg H|B)$ and $i_1 = \text{Pr}(\neg H|\neg E \wedge B)$. It is clearly necessary that $f(1, 0) = f(0, 1) = 0$; for $p = 0$ and

$i = 1$, if H is logically false; and $p = 1$ and $i = 0$ if H is logically true—and in these cases H must be sent to 0, independently of what the data are.

1. Demarcation: $f(1, 0) = f(0, 1) = 0$.

In conjunction with Demarcation, which is violated by r and l ,³ the following is sufficient:

2. Continuity: Any surplus in informativeness succeeds, if the difference in plausibility is small enough.
3. $\forall \varepsilon > 0 \exists \delta_\varepsilon > 0 \forall s_1, s_2, t_1, t_2 \in [0, 1]: s_1 > s_2 + \varepsilon \ \& \ t_1 > t_2 - \delta_\varepsilon \Rightarrow f(s_1, t_1) > f(s_2, t_2)$.

(The s_i are possible values of i , and the t_i are possible values of p .) Indeed, it suffices that Demarcation be conjoined with

3. Continuity in Certainty: Any surplus in informativeness succeeds, if plausibility becomes certainty.
4. $\forall \varepsilon > 0 \forall (t_i)_{i \in N}, (t'_i)_{i \in N} (t_i, t'_i \in [0, 1]): t_i, t'_i \rightarrow i \left\{ \begin{array}{l} 1 \\ 0 \end{array} \right. \exists n \forall m \geq n \forall s_m, s'_m \in [0, 1]: s_m > s'_m + \varepsilon \Rightarrow f(s_m, t_m) > f(s'_m, t'_m)$.

Theorem 2. Let \Pr be a regular probability on \mathcal{L} , let $\{e_i : i \in N\} \subseteq \mathcal{L}$ separate Mod_L , let f be a function of, among others, i and p satisfying Continuity in Certainty and Demarcation, and let \Pr^* be the unique probability measure on the smallest σ -field \mathcal{A} containing the field $\{\text{Mod}(A) : A \in \mathcal{L}\}$ such that for all $A \in \mathcal{L}$: $\Pr(A) = \Pr^*(\text{Mod}(A))$, where $\text{Mod}(A) = \{\omega \in \text{Mod}_L : \omega \models A\}$. Then there exists $X \in \mathcal{A}$ with $\Pr^*(X) = 1$ such that the following holds for every $\omega \in X$, any two contingent $H_1, H_2 \in \mathcal{L}$, and every $H \in \mathcal{L}$:

1. $\omega \models H_1, \omega \not\models H_2 \Rightarrow \exists n \forall m \geq n : f(H_1, E_m^\omega) > 0 > f(H_2, E_m^\omega)$
2. $\omega \models H_1, H_1 \models H_2 \not\models H_1 \Rightarrow \exists n \forall m \geq n : f(H_1, E_m^\omega) > f(H_2, E_m^\omega) > 0$
3. $\omega \not\models H_2, H_1 \models H_2 \not\models H_1 \Rightarrow \exists n \forall m \geq n : 0 > f(H_1, E_m^\omega) > f(H_2, E_m^\omega)$
4. $\models H$ or $\models \neg H \Rightarrow \forall m : f(H, E_m^\omega) = 0$.

3. This defect can be repaired by sticking to the ordinally equivalent r^* and l^* , respectively:

$$r^*(H, E, B) = \lim_{n \rightarrow \infty} \log \left[\frac{\Pr(H|E \wedge B) + 1/n}{\Pr(H|B) + 1/n} \right],$$

$$l^*(H, E, B) = \lim_{n \rightarrow \infty} \log \left[\frac{\Pr(H|E \wedge B) \cdot \Pr(\neg H|B) + 1/n}{\Pr(\neg H|E \wedge B) \cdot \Pr(H|B) + 1/n} \right].$$

However, even Continuity in Certainty is not necessary. The necessary and sufficient condition for revealing the confirmational structure in almost every world when presented separating data is this:

Definition 1. A possibly partial function $f: \mathcal{L} \times \mathcal{L} \times \mathcal{L} \rightarrow \mathfrak{R}$ is a Gaifman and Snir assessment function iff for every probability Pr on a Gaifman and Snir language \mathcal{L} (as described in Section 2) and every $\{e_i : i \in N\} \subseteq \mathcal{L}$ separating Mod_L there is $X \in \mathcal{A}$ with $\text{Pr}^*(X) = 1$ such that for all $\omega \in X$ and all $m \in N$:

- I. $H_1 \models H_2 \not\models H_1, \text{Pr}(H_1|E_m^\omega) \rightarrow m \begin{cases} 1 \\ 0 \end{cases} \Rightarrow \exists n \forall m \geq n : f(H_1, E_m^\omega) > f(H_2, E_m^\omega)$
- II. $\models H_1, \models \neg H_2, \text{Pr}(E_m^\omega) > 0 \Rightarrow f(H_1, E_m^\omega) = f(H_2, E_m^\omega) = 0$.

Condition I, and hence Continuity in Certainty, is violated by r and l . (This defect cannot be repaired by sticking to r^* or l^* .)

Definition 2. Let Pr be a probability on a Gaifman and Snir language \mathcal{L} and let $\{e_i : i \in N\} \subseteq \mathcal{L}$ separate Mod_L . A possibly partial function $f: \mathcal{L} \times \mathcal{L} \times \mathcal{L} \rightarrow \mathfrak{R}$ reveals the confirmational structure of Pr^* -almost every world $\omega \in \text{Mod}_L$ when presented separating $(e_i)_{i \in N}$ iff there is $X \in \mathcal{A}$ with $\text{Pr}^*(X) = 1$ such that for all $\omega \in X$, all contingent $H_1, H_2 \in \mathcal{L}$, and all $H \in \mathcal{L}$:

1. $\omega \models H_1, \omega \not\models H_2 \Rightarrow \exists n \forall m \geq n : f(H_1, E_m^\omega) > 0 > f(H_2, E_m^\omega)$.
2. $\omega \models H_1, H_1 \models H_2 \not\models H_1 \Rightarrow \exists n \forall m \geq n : f(H_1, E_m^\omega) > f(H_2, E_m^\omega) > 0$.
3. $\omega \not\models H_2, H_1 \models H_2 \not\models H_1 \Rightarrow \exists n \forall m \geq n : 0 > f(H_1, E_m^\omega) > f(H_2, E_m^\omega)$.
4. $\models H$ or $\models \neg H \Rightarrow \forall m : f(H, E_m^\omega) = 0$.

f reveals the confirmational structure of almost every world when presented separating data iff for any probability Pr on a Gaifman and Snir language \mathcal{L} and any $\{e_i : i \in N\} \subseteq \mathcal{L}$ separating Mod_L : f reveals the true assessment structure of Pr^* -almost every world $\omega \in \text{Mod}_L$ when presented separating $(e_i)_{i \in N}$.

Theorem 3. A possibly partial function $f: \mathcal{L} \times \mathcal{L} \times \mathcal{L} \rightarrow \mathfrak{R}$ reveals the confirmational structure of almost every world when presented separating data iff f is a Gaifman and Snir assessment function.

One reason why I nevertheless stick to the more general Continuity conditions is that it depends on the underlying convergence theorem which conditions are necessary and sufficient for revealing the true assessment structure in so and so many world when presented such and such data.

More importantly, in the context of theory assessment (Huber, forthcoming) the idea behind the use of these limit considerations is that they provide a theoretical justification for adopting the proposed conditions in the here and now. When assessing several alternative theories we cannot wait until we have arrived at the point of stabilisation for these theories (in fact, we won't know in general when we have passed that point). We need to make our evaluations here and now, where the informativeness and plausibility values are somewhere in between their maximal and minimal values, and we have no idea in which direction they will eventually converge. A theory of theory assessment needs to answer the question what to do when facing such a situation. Continuity does give an answer, but Continuity in Certainty does not. However, we also need to justify this answer—and we do so by appealing to the fact that when we satisfy Continuity in the special case when the plausibility values converge, we almost surely reveal the true assessment structure; and in order to always almost surely reveal the true assessment structure, we always have to be prepared for that convergence to happen, and so we should always satisfy Continuity.

7. Conclusion. I started from the question: Why should one stick to well confirmed theories rather than to any other theories? The answer we got from absolute Bayesian confirmation theory is that one should stick to absolutely well confirmed theories, because absolute confirmation almost surely takes one to true theories. I continued by looking for an answer from incremental Bayesian confirmation theory. This answer should be different from the previous one in order for incremental confirmation to improve on absolute confirmation.

It turned out that three popular measures of incremental confirmation, viz. the distance measure d , the Joyce-Christensen measure s , and the Carnap measure c , give an interesting answer: One should stick to incrementally well confirmed theories, because incremental confirmation almost surely takes one to (the most) informative (among all) true theories.

However, although all measures of incremental confirmation separate contingently true from contingently false theories, not all of them distinguish between informative and uninformative true and false theories. The log-ratio measure r does not distinguish between informative and uninformative false theories, and log-likelihood ratio measure l neither distinguishes between informative and uninformative true nor between informative and uninformative false theories. A sufficient condition for revealing the confirmational structure of almost every world when presented separating data is the conjunction of Continuity and Demarcation, the core principle of the plausibility-informativeness theory of theory assessment (Huber, forthcoming).

REFERENCES

- Bar-Hillel, Yehoshua (1952), "Semantic Information and Its Measures", in Heinz von Foerster, Margaret Mead, and Hans Lukas Teuber (eds.), *Transactions of the Tenth Conference on Cybernetics*. New York: Josiah Macy Jr. Foundation, 33–48.
- (1955), "An Examination of Information Theory", *Philosophy of Science* 22: 86–105.
- Bar-Hillel, Yehoshua, and Rudolf Carnap (1953), "Semantic Information", *British Journal for the Philosophy of Science* 4: 147–157.
- Carnap, Rudolf (1962), *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Carnap, Rudolf, and Yehoshua Bar-Hillel (1952), *An Outline of a Theory of Semantic Information*. Technical Report 247. Cambridge, MA: MIT, Research Laboratory of Electronics.
- Christensen, David (1999), "Measuring Confirmation", *Journal of Philosophy* 96: 437–461.
- Earman, John (1992), *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Fitelson, Branden (1999), "The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity", *Philosophy of Science* 66: S362–S378.
- (2001), "Studies in Bayesian Confirmation Theory." PhD dissertation. Madison: University of Wisconsin.
- Gaifman, Haim, and Marc Snir (1982), "Probabilities over Rich Languages, Testing, and Randomness", *Journal of Symbolic Logic* 47: 495–548.
- Hempel, Carl Gustav (1960), "Inductive Inconsistencies", *Synthese* 12: 439–469.
- (1962), "Deductive-Nomological vs. Statistical Explanation", in Herbert Feigl and Grover Maxwell (eds.), *Minnesota Studies in the Philosophy of Science*, vol. 3. Minneapolis: University of Minnesota Press, 98–169.
- Hilpinen, Risto (1970), "On the Information Provided by Observations", in Jaakko Hintikka and Patrick Suppes (eds.), *Information and Inference*. Dordrecht: Reidel, 97–122.
- Hintikka, Jaakko, and Juhani Pietarinen (1966), "Semantic Information and Inductive Logic", in Jaakko Hintikka and Patrick Suppes (eds.), *Aspects of Inductive Logic*. Amsterdam: North-Holland, 96–112.
- Huber, Franz (forthcoming), "Assessing Theories, Bayes Style", *Synthese*.
- Joyce, James M. (1999), *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Kelly, Kevin T. (1996), *The Logic of Reliable Inquiry*. Oxford: Oxford University Press.
- Levi, Isaac (1961), "Decision Theory and Confirmation", *Journal of Philosophy* 58: 614–625.
- (1963), "Corroboration and Rules of Acceptance", *British Journal for the Philosophy of Science* 13: 307–313.
- (1967), *Gambling with Truth: An Essay on Induction and the Aims of Science*. London: Routledge.
- Milne, Peter (1996), " $\log[P(h|eb)/P(h|b)]$ Is the One True Measure of Confirmation", *Philosophy of Science* 63: 21–26.