Is Evidential Support the Same as Increase-in-Probability?

Tamaz Tokhadze

Author-produced version of an article accepted for publication in *Grazer Philosophische Studien*,(2022).

Abstract

Evidential support is often equated with confirmation, where evidence supports hypothesis H if and only if it increases the probability of H. This paper argues against this received view. As I show, support is a comparative notion in the sense that increase-in-probability is not. A piece of evidence can confirm H, but it can confirm alternatives to H to the same or greater degree; and in such cases, it is at best misleading to conclude that the evidence supports H. I put forward an alternative view that defines support in terms of measures of degree of confirmation. The proposed view is both sufficiently comparative and able to accommodate the increase-in-probability aspect of support. I conclude that the proposed measure-theoretic approach to support provides a superior alternative to the standard confirmatory approach.

Keywords: Support; Confirmation; Evidence; Measures of Confirmation; Measure Sensitivity

1 Introduction

When does a piece of evidence support a hypothesis? As the evidence rarely renders an empirical hypothesis certain, any adequate account of support should be probabilistic, not deductive. A standard positive X-ray test for tuberculosis (TB) does not entail that a patient has TB, but it may still provide good evidence for the presence of the disease.

Following Carnap (1962), epistemologists distinguish two notions of probabilistic support. One is about how firm or probable a hypothesis is on the available evidence. The firmness of a hypothesis can be stated qualitatively: "on the present evidence, it is likely that a patient has the disease" or quantitatively: "the test indicates that the probability of disease in 85%". Support as firmness expresses the static notion of support, which is solely concerned with how probable a hypothesis is on the present evidence.

By contrast, there is also the dynamic notion of support, which is about how a piece of evidence changes or impacts the initial, or prior probability of a hypothesis. So, while the static notion of support is about the *absolute* probability of a hypothesis, the dynamic notion is about the *relationship* between the prior and posterior probability of the hypothesis, relative to some piece of evidence. For instance, when we have the dynamic notion of support in mind, we can see why a positive X-ray test provides evidence for the presence of TB; as the evidence makes the presence of the disease more probable than before.

This paper is solely about the dynamic notion of support. The focus is on finding a general condition (or a set of conditions), which, if satisfied, would license us to conclude that a piece of evidence supports a hypothesis, in the dynamic sense.

The dominant probabilistic view of (dynamic) support appeals to the notion of confirmation, where a piece of evidence, E, supports a hypothesis, H, if and only if E confirms (i.e., increases the probability of) H.[1] In the above diagnostic example, the positive X-ray result is evidence for the presence of TB because it increases the probability of the disease. I call the view that defines the support relation in terms of increase in probability "Support-IP" (IP for Increase in Probability). More fully and precisely:

> Support-IP: E supports H relative to probability measure P iff $P(H|E) > P(H)$.[2]

Support-IP is so ingrained in contemporary formal epistemology that the terms "support" and "confirmation" are often used interchangeably.[3]

---

[1] I will usually omit the modifier "dynamic", as I won't be discussing the static notion of support.

[2] Strictly speaking, confirmation is not a two-term relation involving just E and H, but a three-term relation between E, H, and body of background evidence K. But, in our discussion, this third parameter will not play any role. Therefore, for the sake of readability and simplicity, I will omit K from our notation.

[3] Here are some representative quotes (when necessary, the quoted passages contain

But, certainly, not everybody accepts Support-IP. For instance, Achinstein (1978, 1983, 2000, 2001, 2004, 2013) has been arguing against Support-IP for over 30 years. His main objection is that Support-IP fails to connect the notion of evidence with the notion of acceptance or belief. For Achinstein, if E is evidence for H, then E provides good reason to believe H. And, as I've already explained, E can increase the probability of H without turning H sufficiently probable or believable.

Achinstein's arguments fail to convince most.[4] The common response is that Support-IP is not supposed to guide belief in the first place. After all, epistemologists distinguish two notions of support – static and dynamic. If one is interested in whether a proposition is believable on the evidence, one should appeal to the static notion of support – not the dynamic one; or so the standard reply goes.[5]

some notational changes, for the sake of notational consistency):

> "E confirms (supports, is evidence for) H iff $P(H|E) > P(H)$" (Horwich 1982/2016, 48).
>
> "By confirmation I mean the relation that holds between E and H when E is evidence for H" (Maher 1996, 149).
>
> "E confirms or supports H just in case $P(H|E) > P(H)$" (Howson and Urbach 1993/2006, 92).

[4] For Replies to Achinstein, see Bar-Hillel and Margalit (1979), Maher (1996), and Roush (2004). See Logins (2020) for a detailed discussion and analysis of the debate between Achinstein and his critics.

[5] Achinstein has responded to this response. Briefly, he denies that there are two notions of support: static and dynamic. This paper won't call this standard distinction into question. Instead, and in agreement with most formal epistemologists, I take the distinction for granted.

Another line of criticism against Support-IP comes from the so-called Likelihoodist framework.[6] According to Likelihoodism, support is inherently relational: support is not a two-term relation between evidence and a hypothesis, but a three-term relation between evidence and two competing hypotheses. The Likelihoodist view of support is expressed by the so-called Law of Likelihood (LL) which roughly says that for any two competitor hypotheses H_1 and H_2, E supports H_1 more strongly than H_2 iff E is more likely on the supposition that H_1 than on the supposition that H_2. More precisely:

> LL: For any two competitor hypotheses H_1 and H_2, E supports H_1 over H_2 iff H_1 confers greater probability on E than H_2 does: $P(E|H_1) > P(E|H_2)$.

The Likelihoodist analysis of support also has not attracted many supporters. This is due to several distinct reasons that I cannot go into in this paper. Suffice it to mention two reasons. Firstly, Likelihoodists only offer a theory of comparative support; as LL is restricted to the pairwise comparisons between hypotheses, it cannot, in itself, provide an answer to what a piece of evidence supports simpliciter. And, secondly, Likelihoodists sharply distinguish the question of support from the question of confirmation: E can

[6] The statistician Richard Royall (1997) gave a prominent statement and defence of the Likelihoodist program in statistics. Within epistemology and philosophy of science, Royall's overall approach has been developed by Sober (2008) and Bandyopadhyay et al. (2016).

support H_1 over H_2 (according to LL), but confirm neither of the hypotheses. As most would think that confirmation is at least necessary for support, the analysis of support in terms of LL seems overly weak to capture what a piece of evidence supports.

In this paper, I also argue against Support-IP and propose its alternative. But unlike Achinstein, I do not claim that the dynamic notion of support should guide belief. Instead, and in agreement with Likelihoodists, I argue that support is a comparative notion in the sense that increase-in-probability is not: E can confirm H, but E can confirm some available alternative to H to the same or greater degree. In such situations, it's at best misleading to say that E supports H. However, unlike Likelihoodists, I don't think that the question of support can be independent of confirmation. The alternative to Support-IP that I propose defines support in terms of measures of confirmation. The view – which I call "Support-MDC" (MDC for Maximal Degree of Confirmation) – roughly says the following:

> Support-MDC: E supports H iff the degree to which E confirms H is greater than the degree to which E confirms the alternatives to H.

As I show, Support-MDC is both sufficiently comparative and accommodates the increase-in-probability aspect of Support-IP. My overall conclusion is that Support-MDC is a superior alternative to Support-IP.

The paper runs as follows. The next section discusses what I believe to be, the main shortcoming of Support-IP, which is that the increase-in-probability relation is not sufficiently comparative. Section 3 introduces Support-MDC and provides its three different explications, in terms of three standard measures of degree of confirmation. The superiority of Support-MDC over Support-IP is argued in Section 4. Section 5 discusses and answers, what I take to be, the main worry about Support-MDC: the problem of measure sensitivity. I conclude in Section 6 that Support-MDC should be favoured over Support-IP.

2 Support-IP is not Sufficiently Comparative

In this section, I argue that Support-IP is too weak for capturing what a given body of evidence supports. The problem, as we shall see, is that the increase-in-probability relation is not sufficiently comparative: a piece of evidence, E, can confirm H, but E can also confirm some alternative to H to the same or greater degree.

To illustrate this, let's first consider the following simple diagnostic example:

> *Diagnostic Example*
>
> Suppose you are a physician who wants to determine whether a patient, Eve, has one of the three diseases – d_1, d_2, and d_3. Initially, you consider that the presence of each of the diseases is equally likely. You soon acquire new evidence, denoted by "e", containing various chest X-ray images. Based on the extensive

medical records, you know what type of X-ray image to expect on the condition that Eve has one of these diseases. These medical records suggest the following likelihood distribution: $P(e|d_1) = 0.99, P(e|d_2) = 0.6, P(e|d_3) = 0.03$.

Via Bayes' theorem, the priors and likelihoods enable you to calculate the probability that you are interested in; the posterior probability of a given disease – that is, how probable is the presence of a disease given the data/evidence.[7] The complete probabilistic description of the example is given in the following table:

Priors	Likelihoods	Posteriors
$d_1 = 1/3$	$P(e|d_1) = 0.99$	$P(d_1|e) \approx 0.61$
$d_2 = 1/3$	$P(e|d_2) = 0.6$	$P(d_2|e) \approx 0.37$
$d_3 = 1/3$	$P(e|d_3) = 0.03$	$P(d_3|e) \approx 0.02$

Table 1: Diagnostic Example

As we see, by Bayes' theorem we have $P(d_2|e) \approx 0.37$. Hence, e confirms d_2 in a sense that P $(d_2|e) > P(d_2)$. However, even though e confirms d_2, it seems, at best, misleading to claim that e supports d_2. This is so because there is a hypothesis, d_1, and e supports d_1 much better than d_2. After all,

[7] For instance, by Bayes' theorem $P(d_2|e)$ equals $\frac{P(e|d_2)}{P(e)} * P(d_2)$; and by the law of total probability: $P(e) = P(e|d_1) * P(d_1) + P(e|d_2) * P(d_2) + P(e|d_3) * P(d_3)$; by plugging the numbers: $P(d_2|e) = \frac{0.6}{0.54} * 1/3 \approx 0.37$.

both the posterior probabilities and likelihoods are higher relative to d_1 than relative to d_2.

The example, I believe, illustrates a simple but often overlooked feature of evidential support: that support is a comparative notion. The sole fact that E confirms some hypothesis H does not mean that E supports H, as there might be an available alternative to H that is equally or better confirmed by E. I'll develop a comparative view of support in detail in the next section.

However, before I do this, it will be instructive to compare Diagnostic Example to another, better-known counterexample to Support-IP that has been discussed by Achinstein (2001). As I shall argue, the problem in both cases is similar: Support-IP is unable to capture which one of the two confirmed hypotheses is confirmed to a greater degree by the evidence. I quote Achinstein (ibid., 69) example verbatim:

> *Lottery Example*
>
> b: On Monday all 1000 tickets in a lottery were sold, of which John bought 100 and Bill bought 1. One ticket was drawn at random on Wednesday.
>
> e: On Tuesday all the lottery tickets except those of John and Bill were destroyed, and on Wednesday one of the remaining tickets was drawn at random.

Achinstein asks us to consider the following two hypotheses:

> W_{John}: John won the lottery.

W_{Bill}: Bill won the lottery.

Simple calculations show how the probabilities of W_{John} and W_{Bill} change in light of the evidence:

$P(W_{John}) = \frac{100}{1000} < (W_{John}|e) = \frac{100}{101}$

$P(W_{Bill}) = \frac{1}{1000} < (W_{Bill}|e) = \frac{1}{101}$

As we see, e increases the probability of both W_{John} and W_{Bill}. But Achinstein contends that e supports W_{John} and not W_{Bill}. While I agree with Achinstein on his verdict, I do so for a very different reason. As I've already explained, Achinstein thinks that support should, at least, turn a hypothesis sufficiently probable; as the probability of W_{John} is quite high, he concludes that the evidence supports W_{John} and not W_{Bill}.

Achinstein's diagnosis of the lottery example conflicts with the standard assumption about dynamic support: as we have seen from the TB example, support, in the dynamic sense, does not imply that the supported hypothesis is sufficiently probable. Hence, simply citing the high posterior probability of W_{John} does not imply that the piece of evidence, e, supports W_{John}.

By contrast, I accept the standard assumption about (dynamic) support. For this reason, I don't think that the correct analysis of the lottery example should simply appeal to the high posterior probability of W_{John}. Instead, the problem with Support-IP that the lottery example uncovers is the problem that the diagnostic example uncovers. In both cases, Support-IP fails to capture that one hypothesis is confirmed to a greater degree than its alternative.[8]

To capture the comparativeness of support, I propose to define support in terms of measures of degree of confirmation. This view, which I call Support-MDC (MDC for Maximal Degree of Confirmation), roughly says that evidence supports a proposition iff the evidence confirms it to a greater degree than its alternatives. As we shall see, Support-MDC easily avoids the discussed counterexamples but won't commit us to Achinstein's controversial view that support should always turn a proposition sufficiently probable. This is so because a hypothesis which is confirmed to a greater degree than its alternatives can still have a quite low posterior probability. Hence the account that I offer respects the traditional distinction between the static and dynamic notions of support.

Now, let us proceed to the positive proposal of this paper.

[8] Logins (2020), in response to Achinstein's lottery example, grants that Support-IP is false. Still, he suggests a more restrictive version of Support-IP, which is solely concerned with comparative claims of the following type: "H is more supported given E than it is without E". He proposes and defends the view he calls the *positive probabilistic relevance view about comparative evidential support* (PRCE, for short), which says that "[If] $P(H|E) > P(H)$, then H is more supported given E than it is without E" (ibid., 9).

I don't disagree with Logins. But, this paper is not concerned with such narrow comparative claims about evidence and hypothesis, which is the sole focus of PRCE. Instead, this paper is concerned with dynamic support in general: whether E supports/is evidence for H and not whether H is more supported given E than it is without E. So while my proposed theory of support (Support-MDC) is compatible with Logins' proposal (PRCE), the latter, unlike the former, does not answer the question of the paper: when does a piece of evidence support a hypothesis? For this reason, I won't be concerned with PRCE in this paper.

3 Explicating Support-MDC

I begin by giving the first, rough statement of Support-MDC which will be clarified and explicated shortly.

> Support-MDC: E supports H iff the degree to which E confirms H is greater than the degree to which E confirms the alternatives to H.

Support-MDC involves two unspecified terms: "the degree to which E confirms H" and "alternatives to H". Starting from the latter, I'll specify each of these terms respectively. If support is a comparative notion, then it is not a two-term relation between evidence E and hypothesis H. Instead, support is a three-term relation between E, H, and competitors to H. I call this third relatum "Hyp", which denotes a set of exhaustive and mutually exclusive hypotheses. For simplicity, I assume that Hyp is finite. But, Support-MDC can also be extended to apply to an infinite set of hypotheses (by using a probability density function, instead of a probability mass function). I won't be dealing with a set of infinitely many competitors in this paper.

To make the parameter Hyp explicit, we can restate the definition of Support-MDC as follows:

> Support-MDC: E supports H relative to a set (of mutually exclusive and exhaustive hypotheses) Hyp (with $H \in Hyp$) iff the degree to which E confirms H is maximal with respect to Hyp: i.e., for every x in $Hyp - \{H\}$ (Hyp without H), the degree to

which E confirms H is greater than the degree to which E confirms x.

Now, we also need to define the notion of "the degree to which E confirms H". A common way to define the degree of confirmation relation is to introduce a function, c, that takes a hypothesis and a piece of evidence as inputs and outputs a real number; formally: $c(H, E) = n$, where n is some real number. Plausibly, function c must be related to the probability function, P, in a specific way; so that, if we let t be some threshold value, then function c must satisfy the following conditions:

$c(H, E) > t$ iff $P(H|E) > P(H)$

$c(H, E) = t$ iff $P(H|E) = P(H)$

$c(H, E) < t$ iff $P(H|E) < P(H)$

Trivially, there are infinitely many functions that satisfy the above three conditions.[9] But, only a dozen or so functions have been seriously considered and defended in the literature. Here are some popular, representative measures of confirmation:

Difference: $D(H, E) =_{df} P(H|E) - P(H)$ (Carnap 1962)

[9] These three conditions, in themselves, do not commit to a substantive view about how exactly we ought to measure the degree of confirmation. They only articulate the idea shared by all contemporary confirmation measures; that a function for measuring the degree of confirmation must be associated with some value that demarcates the positive degree of confirmation from the neutral and negative degree of confirmation.

Ratio: $R(H, E) =_{df} \frac{P(E|H)}{P(E)}$ (Milne 1996)

Likelihood-Ratio: $L(H, E) =_{df} \frac{P(E|H)}{P(E|\neg H)}$ (Good 1950; Fitelson 2006)

On measure D, the degree of confirmation between H and E is the difference between the posterior and prior probability of H. Trivially, if $P(H|E) > P(H)$, then the difference between $P(H|E)$ and $P(H)$ is a positive real number. This means that on D the threshold or neutrality value would be 0. For instance, if $D(H, E) > 0$, then $P(H|E) > P(H)$.

On R and L we also have a neutrality value, but, instead of 0, the neutrality value is 1. So, for instance, E confirms H iff $R(H, E) > 1$. (If we wish, we can transform the neutrality value of R and L from 1 to 0, by taking logarithms of these function. But we won't do this, for the sake of readability and simplicity.)

Now, as the notion of degree of confirmation is appropriately specified, we can provide the following schematic explication of Support-MDC, which I call MDC_C.

> MDC_C: E supports H relative to Hyp (with $H \in Hyp$) iff $c(H, E, Hyp)$ is maximal, denoted as "$c_{Max}(E, Hyp) = H$": i.e., for every x in $Hyp - \{H\}$, $c(H, E) > c(x, E)$.

MDC_C is a schematic principle with one unspecified parameter, c; meaning that we get a specific version of Support-MDC once we plug a measure of confirmation in MDC_C, instead of c. For instance, if we plug measure L in MDC_C we get:

MDC_L: E supports H relative to Hyp iff $L_{Max}(E, Hyp) = H$.

But How can we think about the support relation if no hypothesis in Hyp receives a maximal degree of support from the evidence? Firstly, in such a case, we should not categorically assert that the evidence supports some particular hypothesis. For instance, if $c(H_1, E, Hyp) = c(H_2, E, Hyp)$, and if for every x in Hyp where $x \neq H_1 \neq H_2$, $c(H_1, E, Hyp) > c(x, E, Hyp)$, then we should say that E supports H_1 and H_2 equally well relative to their alternatives. In general, giving the complete ordering of competing hypotheses according to some instance of MDC_C would convey *the whole information* about what the evidence supports (relative to a fixed measure of support).

Now, an important thing to note is that Support-MDC is an *ordinal principle*, meaning that for any two measures x and y, if x and y impose the same orderings over a given Hyp, then $MDC_x = MDC_y$. So, the specific numerical value of $c(H, E, Hyp)$ is unimportant. What is important is the ordering that a measure of support imposes on a set of competing hypotheses.

It is well-known that measures D, R, and L are not ordinally equivalent in general; in many cases, they would impose different orderings over a set of competing hypotheses. I'll postpone the discussion of measure sensitivity of Support-MDC and the potential worries that it might raise until Section 5.

First, I'm going to argue that Support-MDC is a superior alternative to Support-IP. I show this by applying Support-MDC to the diagnostic example. Achinstein's example (the lottery example) is more complex and requires a more subtle analysis which I provide in Section 5.

4 Support-MDC is Superior to Support-IP

Support-IP and Support-MDC agree with respect to cases where a piece of evidence confirms just one hypothesis from the set of competing hypotheses. This immediately follows from the feature that all contemporary measures of confirmation have. As we have seen, all measures of confirmation postulate some threshold value t, such that:

$$c(H, E) > t \text{ iff } P(H|E) > P(H)$$

Hence, if H is the only confirmed hypothesis, then $c_{Max}(E, Hyp) = H$.

Now, Support-IP and Support-MDC can give conflicting verdicts where a piece of evidence confirms more than one competing hypotheses. This section argues that, irrespective of choice of measure of confirmation, Support-MDC, unlike Support-IP, provides intuitively correct answers when applied to cases with certain general probabilistic features. We shall say that Support-MDC provides a *robust* verdict when applied to a case iff the verdict is insensitive to choice of measure of confirmation. Hence, what I will show is that, in certain cases, the application of Support-MDC gives correct and robust results.

This can be illustrated by applying Support-MDC to our diagnostic example. Remember, in this diagnostic example we are working with the following probability distribution:

Priors	Likelihoods	Posteriors
$d_1 = 1/3$	$P(e\|d_1) = 0.99$	$P(d_1\|e) = 0.61$
$d_2 = 1/3$	$P(e\|d_2) = 0.6$	$P(d_2\|e) = 0.37$
$d_3 = 1/3$	$P(e\|d_3) = 0.03$	$P(d_3\|e) = 0.02$

Simple calculations show that, the degree to which e supports d_1 is maximal on all considered measures, D, R, and L.[10] More than that, these measures are *ordinally equivalent* with respect to the diagnostic example: that is, they impose the same orderings over Hyp. On all these measures:

$$c(d_1, e) > c(d_2, e) > c(d_3, e)$$

This is in line with the intuitively plausible ordering of these hypotheses. Hence, no matter which of the three measures we choose, Support-MDC, contrary to Support-IP, justifies the intuitively correct verdict that e supports d_1.

[10] As an example, consider measure D:

$$D(d_1, e) = P(e|d_1) - P(d_1) = 0.61 - 1/3 \approx 0.27$$
$$D(d_2, e) = P(e|d_2) - P(d_2) = 0.37 - 1/3 \approx 0.036$$
$$D(d_3, e) = P(e|d_3) - P(d_3) = 0.02 - 1/3 \approx -0.31$$

As we see, d_1 receives the maximal degree of confirmation by e.

Interestingly, there is a certain structural feature that this diagnostic example has, such that, if the probability distribution has that feature, then the evidence would support the same hypothesis according to Support-MDC, relative to all contemporary measures of confirmation (not only relative to the three considered measures).To explain this result, first, we need to introduce the following principle:

> The Weak Law of Likelihood (WLL): E supports H_1 to a greater degree than H_2, if $P(E|H_1) > P(E|H_2)$ and $P(E|\neg H_1) \leq P(E|\neg H_2)$.

Like the Law of Likelihood (LL), WLL is a principle of comparative support: it applies to cases where we want to compare the impact of a piece of evidence on a pair of competing hypotheses. As its name suggests, WLL is strictly logically weaker than LL. LL entails WLL, but not the other way around. We can restate WLL more succinctly by using our measure-theoretic notation:

> WLL: $c(H_1, E) > c(H_2, E)$, if $P(E|H_1) > P(E|H_2)$ and $P(E|\neg H_1) \leq P(E|\neg H_2)$.

On the face of it, WLL seems highly plausible. Both conjuncts of the antecedent of WLL, taken independently, speak in favour of E supporting H_1 over H_2. And if both conjuncts hold simultaneously, this seems to conclusively establish that E supports H_1 over H_2. Here is how Joyce (2003, 8) motivates WLL:[11]

[11] As before, the cited passage contains some notational changes. Instead of WLL,

[WLL] captures one core message of Bayes' theorem for theories of confirmation. Let's say that H_1 is uniformly better than H_2 as predictor of E's truth-value when (a) H_1 predicts E more strongly than H_2 does, and (b) $\neg H_1$ predicts $\neg E$ more strongly than $\neg H_2$ does. According to WLL, hypotheses that are uniformly better predictors of the data are better supported by the data.

WLL certainly enjoys a great deal of intuitive appeal. But what might come as a surprise is that WLL is true relative to all contemporary measures of confirmation (Fitelson 2007, 479). To explain this result in detail, first, let's define the following schematic principle of comparative support, CS_c (CS for Comparative Support and subscript c is an unspecified measure of degree of confirmation):

CS_c: E supports H_1 over H_2 iff $c(H_1, E) > c(H_2, E)$.

When we plug in a specific measure of support instead of parameter c we get a concrete view of comparative support. Now, the surprising fact mentioned above is that, CS_c entails WLL relative to the dozens of measures of support that have been defended in the relevant literature. For instance, if we take our three representative measures, D, R, and L, then WLL can be derived from CS_c, by plugging in one of these measures instead of c (the proofs will be given in the appendix).

Joyce refers to the principle as the "Weak Likelihood Principle". I follow Fitelson (2007) in using the term "WLL".

Given this surprising fact that CS_c entails WLL relative to all contemporary measures of confirmation, we have the following interesting result: if H is a uniformly better predictor of E than any of its alternatives, then E supports H on Support-MDC relative to any contemporary measure of confirmation. More fully and precisely, the result can be stated via the following theorem:

> *The Uniformly Better Predictor Theorem*
>
> Let $H \in Hyp$ and let support be defined as in Support-MDC. Then for every x in $Hyp - \{H\}$:
>
> if $P(E|H) > P(E|x)$ and $P(E|\neg H) \leq P(E|\neg x)$, then E supports H on all contemporary measures of confirmation.

What this theorem shows is that, if a probability model includes a hypothesis, H, and H is the uniformly better predictor of E, then Support-MDC would provide a robust verdict that E supports H. So, no matter which contemporary measure of confirmation we choose, E would still support its uniformly better predictor.[12]

Our diagnostic example includes the uniformly better predictor hypothesis: for every competitors x, $P(E|d_1) > P(E|x)$ and $P(E|\neg d_1) \leq P(E|\neg x)$. Hence, Support-MDC gives a robust verdict when applied to the diagnostic example.

---

[12] I'll only state the proofs relative to three representative measures. The proofs are given in the appendix.

Unfortunately, there are many cases in which Support-MDC does not give robust verdicts. The uniformly better predictor condition does not hold in general; and in such cases, different measures of confirmation can impose different orderings on a set of competitors. For instance, it is well-known that measures D, R, and L are not ordinally equivalent, meaning that they can license opposite comparative judgements in some cases. This makes Support-MDC a measure-sensitive view, where the choice between measures of support can change what a piece of evidence supports.

The next section discusses the measure-sensitivity of Support-MDC and the worry associated with it. I propose and evaluate three strategies for addressing the worry. My overall conclusion will be that measure sensitivity does not pose a serious challenge to Support-MDC.

5 Measure Sensitivity of Support-MDC

There are many measures of confirmation which are ordinally non-equivalent: given the same probability distribution, these measures can impose different orderings on a corresponding set of competing hypotheses. Focusing on our three measures, we can illustrate this by considering the following likelihood distribution over three hypotheses, with an undefined prior probability distribution:

Priors	Likelihoods	Posteriors
$d_1 = x$	$P(e\|d_1) = 0$	$P(d_1\|e) =?$
$d_2 = y$	$P(e\|d_2) = 1$	$P(d_2\|e) =?$
$d_3 = 1 - (x + y)$	$P(e\|d_3) = 0.8$	$P(d_3\|e) =?$

Table 2: Likelihood Distribution

Interestingly, even without knowing the prior distribution over these hypotheses, we can still determine that on measure R, e supports d_2; as $R_{Max}(e, Hyp) = d_2$. Here is how we know this. Table 2 defines the likelihood distribution over the given hypotheses; that is, we know the value of each $P(e|d_i)$. However, in order to determine the value of $R(d_i, e)$ we also need to know the value of $P(e)$; as $R(d_i, e) = \frac{P(e|d_i)}{P(e)}$. But, the value of $P(e)$, in part, depends on the prior distribution, which is not specified in Table 2. Hence $P(e)$ is not a well-defined quantity in this example. However, given measure R, we do not need to know the value of $P(e)$ to impose an ordering on Hyp. This is so because of the following simple theorem:

Simple Theorem

$R(H_1, E) > R(H_2, E)$ iff $P(E|H_1) > P(E|H_2)$ [13]

[13] Here is a simple proof. By definition, for any two propositions H_1 and H_2, we have:

$R(H_1, E) > R(H_2, E)$ iff $\frac{P(E|H_1)}{P(E)} > \frac{P(E|H_2)}{P(E)}$

By algebra (and assuming that $P(E) \neq 0$), $P(E)$ cancels out on the right-hand-side. And we have:

$R(H_1, E) > R(H_2, E)$ iff $P(E|H_1) > P(E|H_2)$

In words: in evaluating inequality $R(H_1, E) > R(H_2, E)$, the expectedness of evidence, $P(E)$, cancels out. Therefore, the inequality, $R(H_1, E) > R(H_2, E)$, is well-defined even without defining the value of $P(E)$.

As this simple theorem makes it explicit, analysing the notion of comparative support in terms of measure R gives us nothing but the familiar Law of Likelihood (LL). Hence, Likelihoodist theory of comparative support is equivalent to analysing the comparative support with Bayesian measure R (Fitelson 2007, 478). Now, given Simple Theorem, R imposes the following ordering on the considered likelihood distribution (as given by Table 2):

$$c(d_2, e) > c(d_3, e) > c(d_1, e)$$

Therefore, on MDC_R, e supports d_2.

By contrast, on measures D and L, the support relation over these hypotheses is undefined. On D and L, if Hyp contains more than two hypotheses, then it is not generally possible to impose an ordering on Hyp without defining the prior distribution first.[14]

The above discussion illustrates that the considered measures can be distinguished by citing the property of *prior-sensitivity*; where R is a prior insensitive measure, while D and L are prior sensitive measures. What this means is that a set of hypotheses, Hyp, can be ordered by R *without defining the prior probability over Hyp*; but the same thing cannot be done via measures L and D, as they are prior sensitive measures of confirmation.

This said, L and D are not prior sensitive in the same way. If we fix the likelihood distribution over Hyp, then L would impose the same ordering as

[14] The sole exception is when all hypothesis except one has the likelihood of 0 on the evidence. But we won't be dealing with such probability models.

R more often than D. Hence, L is less prior sensitive than D. To visualise the degree of prior sensitivity of these measures, let's fix the likelihood distribution as in Table 2, and graph the set of all prior distributions relative to which these measures support either d_2 or d_3. We get the following graph where the blue curve represents measure D and the yellow curve – measure L; and the points (i.e., prior distributions) above a given curve render the judgement that $c(d_2, e) > c(d_3, e)$:

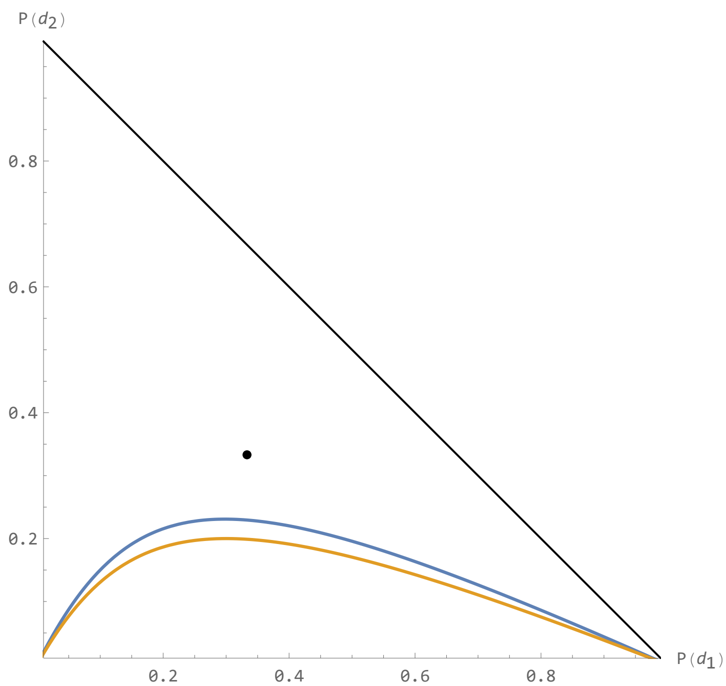

Figure 1: Visualising Prior Sensitivity of D and L

The large black dot in the graph represents the equiprobable distribution; and as we see, on the equiprobable distribution e supports d_2 relative to both measures D and L. However, on some prior distributions, D and L license different judgements. For instance, if we let $P(d_1) = 0.25$ and $P(d_2) = 0.2$, then:

$$L(d_2, e) > L(d_3, e)$$

$$D(d_2, e) < D(d_3, e)$$

And, in general, measure L agrees with prior-insensitive measure R more often than measure D. Hence, L is less prior sensitive than D.

Now, the measure sensitivity of Support-MDC comes into play when we apply Support-MDC to Achinstein's example (the lottery example). To remind the reader, in the lottery example, we have 1000 tickets from which 1 ticket will win. John has 100 tickets, while Bill has just one ticket. Then we get a new piece of evidence, e, that all except John's and Bill's tickets had been destroyed. So, now we know that 1 out of these 101 remaining tickets will win. We have two hypotheses:

W_{John}: John won the lottery.

W_{Bill}: Bill won the lottery.

The new evidence that we received does not change the relative plausibilities of W_{John} and W_{Bill}. Both before and after receiving evidence e, W_{John} is 100 times more likely than W_{Bill}. However, after receiving evidence e, we exclude the possibility that neither John nor Bill won. Hence, W_{John} simply becomes 100 times more likely than $\neg W_{John}$. This means that evidence e does not change the *relative plausibility* of W_{John} vis-à-vis W_{Bill}. However, the *absolute probabilities* of W_{John} and W_{Bill} change significantly. But this changes in absolute probabilities are solely due to difference in prior probabilities of W_{John} and W_{Bill}. As R is a prior insensitive view, using this measure won't give us the verdict that e supports W_{John} over W_{Bill}.[15]

By contrast, as measures D and L are prior sensitive measures, they give us the intuitively correct verdict that e supports only W_{John}. The results of applying these measures to the lottery example are summarised below:

$$D(W_{John}, e) \approx 0.89 > D(W_{Bill}, e) \approx 0.0089$$

$$R(W_{John}, e) \approx 9.9 = R(W_{Bill}, e) \approx 9.9$$

$$L(W_{John}, e) = 900 > L(W_{Bill}, e) = 9.99$$

As we see, Support-MDC agrees with Achinstein's verdict relative to measures D and L. However, on measure R, e is equally good evidence for both hypotheses.

The above discussion illustrates that the choice between R, L, and D can have a significant effect on what a piece of evidence supports. This result gives rise to the following worry:

> If the available evidence is such that it supports different hypotheses relative to different measures, then how should we think

[15] Via Simple Theorem, we saw that analysing comparative support with measure R is logically equivalent to the Law of Likelihood (LL). Therefore, interestingly, Achinstein's lottery example seems to provide a counterexample to LL, if one shares the intuition that the relevant evidence supports W_{John} over W_{Bill}.

about the support relation in such cases?

Now, there are three overall strategies to approach the worry of measure sensitivity. The first strategy is to endorse the uniqueness thesis about measures of confirmation. On this view, there is but *one true measure of confirmation.* And if this kind of uniqueness thesis is true, then the problem of measure sensitivity simply disappears.

While some (e.g., Good 1950; Milne 1996; Fitelson 2006) have defended versions of such uniqueness thesis, most contemporary epistemologists have pluralistic attitudes towards measures of confirmation. This is not a place to go into the debate on the uniqueness thesis about measures of confirmation. Suffice it to say, such uniqueness thesis is an extremely strong and contentious view. So, the appeal to "one true measure of confirmation" is not a very promising strategy for answering the measure sensitivity problem for Support-MDC.

The second strategy is to endorse, what I call, the *qualified pluralism* about measures of confirmation. On this approach, we can grant that no single measure of support is adequate for all cases. Instead, the choice between measures should be made on a *case-by-case basis*: so that, some evidential situations license the use of a unique measure (or a set of ordinally equivalent measures), while some other situations license the use of a different measure. To illustrate the idea behind this qualified pluralism, let's reconsider Achinstein's lottery example. With respect to this lottery case, it seems that one can justifiably favour measure D or L over measure R. Here is why. In the lottery example, the relative probabilities of the competing hypotheses remain the same after we update on the new evidence. Hence, the probabilistic asymmetry between these hypotheses is due to their absolute probabilities (i.e., their prior and posterior probabilities). Now, it seems that, in this specific example, the absolute probabilities are relevant to what the given piece of evidence supports. Hence there is a good reason for favouring measures D and L over measure R, in this particular case. And as D and L impose the same ordering over the relevant competing hypotheses, the choice between them is inconsequential in this case. In making this argument, one can grant that, in some cases, absolute probabilities may be irrelevant or less relevant; still, one may contend that Achinstein's example is not such a case.

Now, the argument for narrowing a set of permissible measures to an ordinally equivalent set of measures might work with respect to Achinstein's lottery example. But, at this point, we have no reason to think that the same strategy would work in general. Hence, such qualified pluralism is not sufficiently motivated to alleviate the worries about measure sensitivity of Support-MDC.

The third and last considered strategy in the face of the measure sensitivity worry is to endorse some form of pluralism about measures of confirmation. On this strategy, we have to grant that, if there is no robust confirmation-theoretic analysis of a given evidential situation, then there is simply no fact of the matter about what this evidence supports. No "absolute" support relation obtains in such cases; we can only make relational claims about what the evidence supports relative to, say, measure R or measure L.

On the supposition that the first two strategies for answering the measure-sensitivity problem (for Support-MDC) are unsuccessful, it seems that pluralism about support is the only available option for us to take. Hence, we have to concede that, in some cases, no “absolute” support relations obtain.

I expect that some readers might not be happy with this conclusion, and look back to Support-IP as a superior alternative to Support-MDC. After all, it seems as if Support-IP is not susceptible to the measure-sensitivity problem: if a hypothesis is confirmed on evidence, then it is confirmed no matter how we measure the degree of confirmation.

But this alleged advantage of Support-IP is illusory. As we’ve already seen, Support-IP is most plausible when the evidence confirms just one hypothesis from the set of competing hypotheses. But such cases are easily handled by Support-MDC as well; relative to all contemporary measures, if evidence confirms just one hypothesis, then the evidence supports this hypothesis relative to any measure of confirmation. In contrast, Support-MDC can deliver robust and intuitively correct results in cases which seem to provide counterexamples to Support-IP (e.g., the diagnostic example). Therefore, even if one is unhappy about accepting a measure-sensitive view of support, the appeal to Support-IP won’t help one to mitigate the worry; as Support-IP cannot even discriminate one confirmed hypothesis from the other. So, instead of addressing the measure sensitivity problem, Support-IP would indiscriminately judge each confirmed hypothesis to be supported by the evidence. This surely cannot be considered a successful dissolution of the measure sensitivity problem.

In summary: even if one is concerned with the measure sensitivity problem, Support-MDC should still be favoured over Support-IP. This concludes the overall argument of the paper.[16]

6 Conclusion

The notion of confirmation is tremendously useful in epistemology. Important epistemic concepts, such as probabilistic relevance, conditional independence and others are defined in terms of confirmation.

However, as I've argued, the relation of support cannot simply be equated with confirmation. This is for the simple reason that support is a comparative notion in the sense that confirmation is not; a piece of evidence can confirm H, but it also can confirm an alternative to H to the same or greater degree. To accommodate the comparativeness of support, I've proposed to define

[16] Besides the problem of measure-sensitivity, one might point out the following potential problem: Support-MDC is sensitive to how we conceptualise the set of competing hypotheses. This is because, on Support-MDC, evidence E can support H relative to one partition, but, E can support some competitor to H relative to a different partition.

While I grant that Support-MDC is partition sensitive in the above sense, I do not consider this to be a shortcoming. After all, Support-IP is also a partition sensitive view. For instance, consider the diagnostic example. If we exclude hypothesis d_3 from the outset, then the available evidence would no longer confirm d_2. Hence, the partition sensitivity is not an idiosyncratic feature of Support-MDC.

support in terms of measures of confirmation.

This proposed view, Support-MDC, can easily handle some of the counterexamples against the increase-in-confirmation account of support (Support-IP). However, there are cases where Support-MDC is not as easy to apply: the cases where choice of measure of confirmation becomes important.

While I do not think that any adequate theory of support can completely escape the problem of measure sensitivity, the project of articulating a more robust theory of support (compared to Support-MDC) is very well motivated. Whether this project succeeds or not, the conclusion of the paper would still be the same: Support-MDC is a superior alternative to Support-IP.

It is also worth noting that Support-MDC has an important practical advantage over Support-IP: the former, unlike the latter, helps us to avoid overestimating the epistemic significance of confirmation. As we have seen, a piece of evidence can confirm H, but the same evidence might not support or "speak in favour" of H. For this reason, one should not attach too much weight to the increase-in-probability relation, unless one has evaluated the evidence with respect to a set of relevant competing hypotheses.

Acknowledgements

I am very grateful to Brett Topey, Prasanta Bandyopadhyay and an anonymous reviewer for this journal for written feedback and suggestions. Many thanks to Tony Booth, Corine Besson, and Nina Abesadze for helpful discussions on various parts and aspects of this paper. Thanks a lot to Simon Mcgregor for helping me with a part of the proof of the theorem in the appendix.

References

Achinstein, Peter 1978. "Concepts of Evidence." *Mind* 87(345), 22–45.

Achinstein, Peter 1983. *The Concept of Evidence*. Oxford: Oxford University Press.

Achinstein, Peter 2000. "Why Philosophical Theories of Evidence Are (And Ought to Be) Ignored by Scientists." *Philosophy of Science* 67, 180–192.

Achinstein, Peter 2001. *The Book of Evidence*. Oxford: Oxford University Press.

Achinstein, Peter 2004. "A Challenge to Positive Relevance Theorists: Reply to Roush." *Philosophy of Science* 71(4), 521–524.

Achinstein, Peter 2013. *Evidence and Method: Scientific Strategies of Isaac Newton and James Clerk Maxwell*. Oxford: Oxford University Press.

Bandyopadhyay, Prasanta S, Gordon G Brittan, and Mark L Taper 2016. *Belief, Evidence, and Uncertainty: Problems of Epistemic Inference*. New York: Springer.

Bar-Hillel, Maya and Avishai Margalit 1979. "In Defence of the Classical Notion of Evidence." *Mind* 88, 576–583.

Carnap, Rudolf 1962. *Logical Foundations of Probability*. 2nd. Chicago: University of Chicago Press.

Fitelson, Branden 2006. "Logical Foundations of Evidential Support." *Philosophy of Science* 73(5), 500–512. DOI: 10.1086/518320.

Fitelson, Branden 2007. "Likelihoodism, Bayesianism, and Relational Confirmation." *Synthese* 156, 473–489. DOI: 10.1007/s11229-006-9134-9.

Good, Irving John 1950. *Probability and the Weighing of Evidence*. London: Griffin.

Horwich, Paul 1982/2016. *Probability and Evidence*. Cambridge: Cambridge University Press.

Howson, Colin and Peter Urbach 1993/2006. *Scientific Reasoning: The Bayesian Approach*. 3rd. La Salle, IL: Open Court.

Joyce, James 2003. "Bayes' Theorem." *Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), Edward N. Zalta (ed.), URL: https://plato.stanford.edu/archives/fall2021/entries/bayes-theorem/ (visited on 04/01/2022).

Logins, Artūrs 2020. "Is an Increase in Probability Always an Increase in Evidential Support?" *Erkenntnis*, 1–25. DOI: 10.1007/s10670-020-00241-4.

Maher, Patrick 1996. "Subjective and Objective Confirmation." *Philosophy of Science* 63(2), 149–174.

Milne, Peter 1996. "log [p (h/eb)/p (h/b)] is the One True Measure of Confirmation." *Philosophy of Science* 63, 21–26.

Roush, Sherrilyn 2004. "Discussion Note: Positive Relevance Defended." *Philosophy of Science* 71(1), 110–116.

Royall, Richard 1997. *Statistical Evidence: A Likelihood Paradigm.* New York: Chapman Hall.

Sober, Elliott 2008. *Evidence and Evolution: The Logic Behind The Science.* New York: Cambridge University Press.

7 Appendix

In this appendix we will prove the Uniformly Better Predictor Theorem by using measures D, R, and L. The theorem states the following:

> *The Uniformly Better Predictor Theorem*
>
> Let $H \in Hyp$ and let support be defined as in Support-MDC. Then for every x in $Hyp - \{H\}$:
>
> if $P(E|H) > P(E|x)$ and $P(E|\neg H) \leq P(E|\neg x)$, then E supports H on measures D, R, and L.

Now, according to Support-MDC, E supports H relative to Hyp iff the degree to which E confirms H is maximal (denoted as $c_{Max}(E, Hyp) = H$). So, to prove this theorem for some measure of confirmation c, we need to show that:

> if $P(E|H) > P(E|x)$ and $P(E|\neg H) \leq P(E|\neg x)$, then $c_{Max}(E, Hyp) = H$

This is exactly what we are going to demonstrate with respect to measures D, R, and L.

The theorem is trivially true for measures R and L. To simplify the exposition, let X be some arbitrary member of Hyp, where $H \neq X$. And let A denote $P(E|H) > P(E|X)$ and B denote $P(E|\neg H) \leq P(E|\neg X)$. Now, by logic we have:

$$(A \wedge B) \Rightarrow A \tag{1}$$

And by definition:

$$R(H, E) > R(X, E) \iff A \tag{2}$$

As X can be any competitor to H, we conclude that for every competitors x:

$$P(E|H) > P(E|x) \text{ and } P(E|\neg H) \leq P(E|\neg x) \Rightarrow R_{Max}(E, Hyp) = H \tag{3}$$

Hence, the proof of the theorem with respect to measure R.

Regarding measure L: it is trivial that:

$$(A \wedge B) \Rightarrow L(H, E) > L(X, E) \tag{4}$$

And given that X can be any competitor to H, (4) is already equivalent to:

$$P(E|H) > P(E|x) \text{ and } P(E|\neg H) \leq P(E|\neg x) \Rightarrow L_{Max}(E, Hyp) = H \tag{5}$$

Hence, the proof of the theorem with respect to measure L.

For measure D, the theorem is a bit more cumbersome to prove. As before, we assume that all relevant conditional probabilities are well-defined; so, for the relevant hypotheses H and X and evidence E, the following quantities are non-zero: $P(H)$, $P(X)$, $P(\neg H)$, $P(\neg X)$, $P(E)$. Now we will prove that, if $P(E|H) > P(E|X)$ and $P(E|\neg H) \leq P(E|\neg X)$, then $P(H|E) - P(H) > P(X|E) - P(X)$; so, if the conditions of the theorem are satisfied, then $D(H, E) > D(X, E)$. Here is one way to prove it.

By the axioms of probability (and the definition of conditional probability), if we multiply $P(E|H) > P(E|X)$ by $P(H)P(X)$ we get:

$$P(H \wedge E)P(X) > P(X \wedge E)P(H) \tag{6}$$

Similarly, multiply $P(E|\neg H) \leq P(E|\neg X)$ by $P(\neg H)P(\neg X)$ to get:

$$P(\neg H \wedge E)P(\neg X) \leq P(\neg X \wedge E)P(\neg H) \tag{7}$$

Via the axioms of probability, remove complements (e.g., substitute $P(\neg H \wedge E)$ with $P(E) - P(H \wedge E)$):

$$(P(E) - P(H \wedge E))(1 - P(X)) \leq (P(E) - P(X \wedge E))(1 - P(H)) \tag{8}$$

Multiply out and rearrange equation (8):

$$P(H \wedge E)P(X) \leq P(H \wedge E) - P(X \wedge E) + P(E)(P(X) - P(H)) + P(X \wedge E)P(H) \tag{9}$$

Now, given equations (6) and (9), we have:

$$P(H \wedge E) - P(X \wedge E) + P(E)(P(X) - P(H)) + P(X \wedge E)P(H) \geq P(H \wedge E)P(X) > P(X \wedge E)P(H) \tag{10}$$

And subtract $P(X \wedge E)P(H)$ from equation (10):

$$P(H \wedge E) - P(X \wedge E) + P(E)(P(X) - P(H)) \geq P(H \wedge E)P(X) - P(X \wedge E)P(H) > 0 \tag{11}$$

From equation (11):

$$P(H \wedge E) - P(X \wedge E) + P(E)P(X) - P(E)P(H) > 0 \tag{12}$$

By algebra:

$$P(H \wedge E) - P(E)P(H) > P(X \wedge E) - P(E)P(X) \tag{13}$$

Finally, divide equation (13) by $P(E)$ (and via the definition of conditional probability):

$$P(H|E) - P(H) > P(X|E) - P(X) \tag{14}$$

Hence, the theorem is proved with respect to measure D as well.