A HYPERINTENSIONAL CRITERION OF IRRELEVANCE

STEPHAN KRÄMER

ABSTRACT. On one important notion of irrelevance, evidence that is irrelevant in an inquiry may rationally be discarded, and attempts to obtain evidence amount to a waste of resources if they are directed at irrelevant evidence. The familiar Bayesian criterion of irrelevance, whatever its merits, is not adequate with respect to this notion. I show that a modification of the criterion due to Ken Gemes, though a significant improvement, still has highly implausible consequences. To make progress, I argue, we need to adopt a hyperintensional conception of content. I go on to formulate a better, hyperintensional criterion of irrelevance, drawing heavily on the framework of the truthmaker conception of propositions as recently developed by Kit Fine.

This is the penultimate version of a paper accepted for publication in Synthese.
Do not cite without permission.

1. INTRODUCTION

In the context of any inquiry, we need to distinguish between evidence that is relevant to the problem at hand, and evidence that is irrelevant. The distinction is of some importance. Discarding relevant evidence as irrelevant increases the likelihood of error and bad decision-making, and treating irrelevant evidence as relevant results at best in a waste of resources. In view of its centrality to our cognitive and practical lives, the notion of relevance that is in play here carries great philosophical interest, and an adequate explication of the notion would be highly desirable.[1]

---

[1] As Cohen has emphasized (cf. 1994: p. 171f), there are also important applications of the notion of relevance which cannot readily be represented as concerning a relation of evidence to hypothesis. In this paper, I restrict attention to applications which can naturally be so understood. I suspect, however, that the tools I shall employ in accounting for these instances of relevance can also fruitfully be applied in a much greater range of cases. I hope to further pursue this question in future work.

At first glance, it may appear that such an explication is afforded by the understanding of relevance in terms of a change in probabilities. Roughly speaking, a piece of evidence is relevant to a given hypothesis just in case it makes the hypothesis either more or less probable than it would otherwise be. This suggestion is given a precise formulation in the shape of the usual Bayesian standard for irrelevance:[2]

> **(BI):** A piece of evidence E is irrelevant to a hypothesis H iff $\Pr(H|E) = \Pr(H)$.

Unfortunately, as Ken Gemes (2007) has convincingly argued, this explication of our notion is not satisfactory.[3] In a rough approximation, the problem is that evidence may bear on a hypothesis both in a positive and a negative way, so that its probabilistic effects cancel each other out.[4] In such a case, (BI) issues a verdict of irrelevance. But in real situations of inquiry and decision-making, Gemes points out, we would not, and we should not, discard such pieces of evidence as irrelevant (cf. 2007: p. 162f).

Gemes then goes on to propose a new account of irrelevance which avoids the problematic feature of (BI). The account has two main components. The first is a modified criterion of Gemesian irrelevance that replaces (BI):[5]

> **(GI):** A piece of evidence E is irrelevant to a hypothesis H iff for every part E′ of E and every part H′ of H, $\Pr(H'|E') = \Pr(H')$.

In contrast to (BI), (GI) invokes a notion of a *part* of a content or proposition, such as a piece of evidence or a hypothesis.[6] Although we have some intuitive grip on that notion,

---

[2] Pr(H) denotes the (prior) probability of the hypothesis H, and Pr(H|E) denotes the (posterior) probability of H given the evidence E. The standard provided by (BI), and any of its refinements to be considered below, is thus relative to a suitable prior probability distribution. Following common practice in the debate, we won't worry here about the exact nature of the probabilities in question. The notion of irrelevance is often further relativized to a body of background information K and is then taken to be characterized by the condition that $\Pr(H|E \wedge K) = \Pr(H|K)$. I have left out reference to K throughout. Doing so facilitates comparison with Gemes' 2007 account to be considered below, and the difference is of no import for our purposes.

[3] This is not to say that (BI) may not amount to an adequate explication of *a* useful notion of irrelevance. The claim is that the notion sketched above is not adequately explicated by (BI).

[4] Cf. (Gemes, 2007: p. 162). As Gemes highlights, this kind of complaint against (BI)'s identification of relevance with probabilistic relevance is not new; it was already made in 1929 by John Maynard Keynes (1929: p. 79).

[5] See (Gemes, 2007: p. 165). Gemes' formulation has 'content part' in place of 'part', which is Gemes' term for his explication of the notion of a part of a content. For our purposes, it is better to view as separate the specific account of that notion and the proposed revision of (BI). The term 'part', here and throughout, is to be understood as 'proper or improper part'.

[6] I use the terms 'content' and 'proposition' interchangeably. I have found it convenient to speak of evidence and hypotheses as themselves propositions rather than sentences expressing these propositions.

it clearly stands in need of explication. The second component of Gemes' proposal is accordingly a precise account of this notion, which Gemes has developed and defended in more detail in earlier papers (1994; 1997).

I have two primary aims in this paper, one destructive, the other constructive. The destructive aim is to reveal some problems both for the criterion (GI) and for Gemes' specific account of content parts. The constructive aim is to formulate a better account of relevance. This may be done, I argue, by firstly, replacing the Gemesian account of content parts with the rival account recently developed by Kit Fine (2015), and secondly, by tweaking (GI) somewhat. Both Fine's account of partial content[7] and the tweaks to (GI) that I propose draw heavily on the hyperintensional distinctions made available in Fine's *truthmaker* account of content. A secondary aim of the paper is therefore to bring out, on the one hand, the importance and fruitfulness of the recognition of relevance as a hyperintensional relation, and on the other hand, the usefulness of the truthmaker account as a tool for explicating and studying hyperintensional concepts.

The plan for the paper is as follows. In section 2, I present an example in which Gemes' account yields a counter-intuitive verdict of irrelevance. In this case, I argue, the problem lies with Gemes' account of content parts. On an intuitive understanding of that notion, the cases are consistent with (GI). Section 3 argues that the difficulties arising for the Gemesian view may be traced to his insistence that logically equivalent sentences must have the same content. I show that we obtain a much more plausible version of (GI) once we replace Gemes' intensional account of content parts with its hyperintensional Finean rival, which allows for logically equivalent sentences to have different contents, and hence content parts. Section 4 argues that intuitively, even under the Finean interpretation of 'part', (GI) is subject to counter-examples. To avoid them, two changes are proposed. In a first step, we broaden our attention to consider not just parts of the hypothesis, but parts of any disjunct of the hypothesis. In a second step, we restrict attention to what I call *helpful* parts. Roughly, these are parts whose truth brings us closer to the truth of the hypothesis than we are independently of them. Section 5 concludes.

2. A counter-example

I shall begin by briefly reviewing the case that Gemes uses to argue against (BI) and to motivate his own account (cf. 2007: p. 162f). Suppose two dice A and B have been tossed, and consider the following pair of evidence and hypothesis:

---

[7]   I use the phrases 'content part' and 'partial content' interchangeably. The former is favoured by Gemes, the latter by Fine.

(E_1): A came up 1, and B came up either 1, 3, 5, or 6.

(H_1): A came up odd, and B came up even.

In this case (given natural background assumptions), $\Pr(H_1|E_1) = \Pr(H_1)$.[8] According to (BI), therefore, the evidence E_1 is irrelevant to H_1. But this seems to be the wrong result. It would appear irrational, in the context of an inquiry into H_1, to discard the information that E_1 holds as irrelevant. Gemes' example therefore constitutes a strong case against (BI) as an explication of the target notion of irrelevance.

The case also lends some support to Gemes' own account, which implies that E_1 is relevant to H_1. Clearly, the proposition

(E_1'): A came up 1.

is probabilistically relevant to H_1, for $\Pr(H_1|E_1') = 1/2 \neq 1/4 = \Pr(H_1)$. Intuitively, E_1' would certainly seem to qualify as a part of E_1, in which case (GI) implies that E_1 is relevant to H_1. Similarly, E_1 itself is probabilistically relevant to the hypothesis

(H_1'): A came up odd.

For $\Pr(H_1'|E_1) = 1 \neq 1/2 = \Pr(H_1')$. Intuitively, H_1' certainly qualifies as a part of H_1, which again would mean that (GI) counts E_1 relevant to H_1. As Gemes shows, the relevant claims of parthood are vindicated on his account of content parts. So it seems that Gemes can handle the example quite convincingly.[9]

Unfortunately, there are many other cases which intuitively, as far as relevance is concerned, are of exactly the same sort as the previous example, but nevertheless receive the opposite verdict on Gemes' account. Here is one such case. Suppose two fair coins A and B have been tossed, and consider the following pair of evidence and hypothesis:

(E_2): A came up heads.

(H_2): (A came up heads or B came up tails) and (A came up tails or B came up heads).

---

8   The obvious prior probability of H_1 is 1/4, which is also the probability of H_1 given E_1, since E_1 is compatible with four equally probable outcomes of the tosses, exactly one of which makes H_1 true.

9   Cf. (Gemes, 2007: pp. 165f). Note that to deal with this particular example, a more modest deviation from (BI) than is embodied in (GI) would have been sufficient. For as we saw, (GI) overdetermines, as it were, the result that the E_1 is relevant to H_1 in that we have both that part of the evidence is probabilistically relevant to the hypothesis as a whole, and that the evidence as a whole is probabilistically relevant to part of the hypothesis. But this feature is specific to the current example. The next example in the main text plausibly can only be captured by allowing as sufficient for relevance the probabilistic relevance of the evidence for part of the hypothesis. For an example of the 'converse' sort, consider any hypothesis H and the corresponding evidence that Bill said H was false and Bob said H was true, where Bill and Bob are generally reliable sources equally likely to be wrong (or lying) with respect to H or any part of H.

It is easy to see that $\Pr(H_2) = \Pr(H_2|E_2) = 1/2$. The outcomes that render H_2 true are exactly the both-heads and the both-tails outcomes, and the probability that one of these two obtains is $1/2$ independently of the evidence that A came up heads as well as given that evidence. So according to (BI), E_2 is irrelevant to H_2. But just as in Gemes' case, this result is implausible. Moreover, on an intuitive understanding of content parts, the case also seems to fit with (GI). Clearly, E_2 is probabilistically relevant to the hypothesis

(H$_2'$): A came up heads or B came up tails.

For $\Pr(H_2'|E_2) = 1 \neq 3/4 = \Pr(H_2')$. Intuitively, H_2' certainly seems to qualify as part of H_2, in which case (GI) yields the desirable result that E_2 is relevant to H_2.

However, on Gemes' account, neither H_2' nor H_2's other conjunct that A came up tails or B came up heads is a content part of H_2. On this account, for the content of a sentence α to be part of the content of a sentence β is for α to be a special kind of logical consequence of β. Gemes offers two characterizations of the additional condition that has to be satisfied, one relatively informal, the other more formal. Suppose that α is a logical consequence of β. Then on the informal version, in order for α's content to be a Gemesian part of the content of β, it must be the case that: there is no logical consequence γ of β such that γ logically entails, but is not logically entailed by, α, and 'all the vocabulary of [γ] occurs (essentially) in α' (Gemes, 2007: p. 164).[10]

More formally characterized, the condition a logical consequence α of β must satisfy in order for α's content to qualify as as Gemesian part of the content of β is that every relevant model of α can be extended to a relevant model of β:

(GP): The content of α is part of the content of β iff $\beta \vdash \alpha$, and every relevant model of α can be extended to a relevant model of β.

---

[10] The quoted phrase, and in particular the parenthetical qualification, is the bit which requires further clarification. Firstly, the vocabulary of γ in question is supposed to be non-logical vocabulary, otherwise the fact that $P \wedge Q$ is a stronger consequence of $P \wedge Q$ than $P \vee Q$ would not prevent the latter from being a part of the $P \wedge Q$, as Gemes clearly intends it to do (cf. 1994: p. 603). Secondly, from the official definition offered in (Gemes, 1994: p. 605), we may extract that a piece of non-logical vocabulary is said to occur essentially in a sentence just in case there is no logically equivalent sentence in which it does not occur. The motivation for the restriction to α's essential vocabulary is to ensure that logically equivalent sentences stand in the same parthood relations (cf. (Gemes, 1994: p. 604f)). To see the point, note that without the restriction, P is part of $P \wedge Q$, but the logically equivalent $P \wedge (Q \vee \neg Q)$ is not, since $P \wedge Q$ itself is a stronger logical consequence of $P \wedge Q$ which contains all the vocabulary in $P \wedge (Q \vee \neg Q)$.

Here, a relevant model is one that assigns truth-values only to the relevant sentence letters in a formula, and a sentence letter is relevant in a formula iff changing its truth-value within a given model can change the truth-value of the formula (cf. ibid). Thus, 'P' is relevant in '$P \lor Q$' and in '$\neg P \land Q$', but irrelevant in both '$P \lor \neg P$' and '$Q \lor (P \land Q)$'.

To apply the account to our example, we first formalize evidence and hypothesis within a propositional language. Let 'P' and 'Q' stand for the propositions that coin A came up heads and that coin B came up heads, respectively. Our hypothesis H_2 may then be written as '$(P \lor \neg Q) \land (\neg P \lor Q)$'. We now see that '$P \lor \neg Q$' has a relevant model – the model assigning True to 'P' and False to 'Q' – which cannot be extended to a relevant model of '$(P \lor \neg Q) \land (\neg P \lor Q)$'. And likewise '$\neg P \lor Q$' has a relevant model – the model assigning False to 'P' and True to 'Q' – which also cannot be extended to a relevant model of '$(P \lor \neg Q) \land (\neg P \lor Q)$'. Nor are there other parts of the content of '$(P \lor \neg Q) \land (\neg P \lor Q)$' to which the evidence P might be probabilistically relevant. For in order for the content of α to be part of the content of '$(P \lor \neg Q) \land (\neg P \lor Q)$', α must not contain any sentence letters except for 'P' and 'Q'. If α contains only 'P' or only 'Q', then in order for it to be a logical consequence of '$(P \lor \neg Q) \land (\neg P \lor Q)$', it must be logically true, so the evidence cannot be probabilistically relevant to it. If α contains both 'P' and 'Q', then α must be logically equivalent to '$(P \lor \neg Q) \land (\neg P \lor Q)$', otherwise it will have a relevant model that is not, and cannot be extended to, a relevant model for '$(P \lor \neg Q) \land (\neg P \lor Q)$'.[11]

3. Partial content

At first glance, it seems an odd feature of (GP) that the conjuncts of a conjunction may fail to be parts of the conjunction. The relation between a conjunct P and a corresponding conjunction $P \land Q$ seems to be the very paradigm of the relation of content part; it is no accident that the content parts of evidence and hypothesis figuring in the example by which Gemes seeks to motivate (GI) are conjuncts of evidence and hypothesis, respectively. So one may wonder why Gemes chooses to restrict the relation of content part in a way that rules out some instances of the conjunct-conjunction relation. The reason is

---

[11] We can construct a similar case pertaining to parts of the evidence. Consider some hypothesis H with prior probability 1/2 and two generally reliable sources Bill and Bob who are equally likely to be wrong (or lying) about H, and let the evidence be that (Bob or Bill said H is true) and (Bob or Bill said H is false). Let it be given that both Bob and Bill said either that H is true or that H is false, so that the evidence may be represented as $(P \lor Q) \land (\neg P \lor \neg Q)$. The evidence as a whole is intuitively, but not probabilistically, relevant for the hypothesis, whereas the conjuncts are probabilistically relevant. For the same reasons as before, however, the conjuncts do not qualify as parts of the evidence on Gemes' account.

that he is forced to do so given two other principles he wishes to uphold. I explain them in turn.

As Gemes highlights, the most important feature, for his purposes, of the notion of content part is that not every logical consequence of P, and in particular not every disjunction $P \lor Q$, counts as a part of P. For if it did, (GI) would yield an almost empty notion of irrelevance. The reason is that in almost all cases of a piece of evidence E and a hypothesis H, the disjunction $E \lor H$ is probabilistically relevant to H. Hence if $E \lor H$ is a part of E, by (GI), E automatically comes out relevant to H. It is intuitively quite plausible to deny that $P \lor Q$ is always a part of P. To vary a point made by Gemes, we should otherwise have to say for arbitrary P and Q that the evidence Q conclusively confirms part of the hypothesis P, which seems bizarre. Similarly, we are in no way tempted to count P as partially true, purely on the strength of the truth of Q, and hence $P \lor Q$. (Cf. (Gemes, 2007: p. 164); (Gemes, 1994: p. 597ff); see also Fine (2013).)

The second principle is a principle of *intensionality*. It says that pairs of logically equivalent sentences always have the same content, and thus the same content parts.[12] Given this principle, the claim that every conjunction $P \land Q$ contains its conjuncts as parts implies that an arbitrary logical consequence Q of a given proposition P is always also a part of P, in contradiction of the first principle. For suppose Q is a logical consequence of P. Then P is logically equivalent to $P \land Q$. But by the conjunction principle, Q is part of $P \land Q$, and hence of P.

Given these two principles, then, Gemes has no choice but to restrict the principle that conjuncts of a conjunction are parts of the conjunction. As we have seen, the restriction he chooses yields some counter-intuitive denials of parthood, and in conjunction with (GI), some counter-intuitive denials of relevance. We should note, moreover, that the denials of parthood cannot be motivated by the sorts of considerations that Gemes uses to rule out arbitrary disjunctions as parts. For intuitively, there is no problem with saying that $(P \lor \neg Q) \land (\neg P \lor Q)$ is partially true, or that a part of it has been conclusively confirmed, on the strength of $P \lor \neg Q$.

Indeed, these kinds of intuitive considerations seem to speak heavily in favour of accepting the conjunction principle and instead giving up on the intensionality principle. For intuitively, there is also no problem with saying that $P \land (P \lor Q)$ is partially true, or that a part of it has been conclusively confirmed, on the strength of Q. At the same time, there is no temptation at all to infer from this that the logically equivalent statement P is partially true, or that part of P has been conclusively confirmed.

---

[12] This principle is not discussed in Gemes (2007), but it plays a central role in Gemes' development of his view in (Gemes, 1994: cf. esp. pp. 601-605).

All this suggests that a hyperintensional conception of content may be better suited for the explication of the notion of content part, and thus for giving a plausible interpretation to the irrelevance criterion (GI). To confirm this conjecture, I will now briefly sketch the *truthmaker* conception of content recently developed by Kit Fine and explain how Fine proposes to explicate the notion of parthood within his framework.[13] I will then apply Fine's view to our case and show that it gives the desired results.

The best way to introduce Fine's truthmaker conception of content is to contrast it with the familiar possible worlds conception of content. On the latter view, a proposition is identified with the set of possible worlds at which the proposition is true. On the truthmaker view, a proposition is instead described in terms of the set of *states* which make the proposition true.[14] Roughly speaking, states are like possible worlds except in that they need not be possible, and they need not be (complete) worlds. That is, whereas a world is complete in the sense that it settles the truth-value of every proposition, a state may be incomplete and leave open the truth-values of many propositions.

States are assumed to have mereological structure. In particular, given any states $s, t, u, \ldots$ we may form their *fusion* $s \sqcup t \sqcup u \sqcup \ldots$ which is the smallest state containing all of $s, t, u, \ldots$ as parts. We shall take for granted two important principles about propositions and the states that make them true, or *verify* them:

(Closure): If a proposition P is verified by each of some states $s, t, u, \ldots$ then it is also verified by their fusion $s \sqcup t \sqcup u \sqcup \ldots$

(Convexity): If a proposition P is verified by each of the states s and u, then P is also verified by any state t which is both part of u and has s as a part.

It is important to note, however, that *verification* is not assumed to be *monotonic* in the sense that if a proposition is verified by a state s, then it is also verified by any bigger state s' of which s is a part. Roughly speaking, on Fine's construal of verification, in order for a state to verify a proposition, every part of the state must play a part, must be actively involved, as it were, in the verification of the proposition.[15]

---

[13] The fullest published exposition by Fine of the truthmaker conception of content and the notion of partial content is given in Fine (2015) in the context of a discussion of Angell's logic of analytic entailment. A more general presentation and discussion of the framework is contained in the as yet unpublished manuscripts Fine (msa) and Fine (msb).

[14] While this may make it sound as though the views make incompatible claims about the same kind of thing, viz. propositions, it is not necessary for my purposes that we think of the views in this way. We may instead take them to concern different concepts of propositions, suited to different theoretical purposes.

[15] Strictly speaking, Fine distinguishes a number of different conceptions of verification. I am here concerned with what Fine calls exact verification, which is the basic notion of verification in terms of which he defines other, looser conceptions. Cf. (Fine, 2015: pp. 7f, 20f), and (Fine, msa: p. 35f). Fine

On the possible worlds picture, by fixing which worlds make a proposition P true, we ipso facto also fix which worlds make P false, namely all worlds which do not make P true. We are thereby in a position to say which worlds make the negation $\neg P$ true, namely those worlds which make P false. On the truthmaker picture, this is not so. A state which does not verify a proposition P need not therefore falsify it, and correspondingly, it need not verify $\neg P$. As a result, we have to separately specify the verifiers and the falsifiers of a given proposition.[16] A proposition P is therefore identified with an ordered pair $\langle P^+, P^- \rangle$ of a non-empty set P^+ of states verifying P and a non-empty set P^- of states falsifying P. Operations of conjunction, disjunction and negation on the propositions are then defined as follows (for X a set of states, we write X° for the smallest closed and convex set containing X):

$$
\begin{aligned}
(\neg P)^+ &= P^- \\
(\neg P)^- &= P^+ \\
(P \wedge Q)^+ &= \{s \sqcup t : s \in P^+ \text{ and } t \in Q^+\}^\circ \\
(P \wedge Q)^- &= (P^- \cup Q^-)^\circ \\
(P \vee Q)^+ &= (P^+ \cup Q^+)^\circ \\
(P \vee Q)^- &= \{s \sqcup t : s \in P^- \text{ and } t \in Q^-\}^\circ
\end{aligned}
$$

Fine now proposes the following account of what it is for a proposition P to be part of a proposition Q (cf. (Fine, 2015: p. 8ff, 19)):

(FP) P is a part of Q iff (i) every verifier of P is part of a verifier of Q
                                  (ii) every verifier of Q has a part that verifies P
                                  (iii) every falsifier of P is a falsifier of Q

Like (GP), (FP) has the desirable consequence that $P \vee Q$ is not in general a part of P, since a verifier of Q will not in general be part of a verifier of P. (Note that for this result it is important that verification is non-monotonic.) Unlike (GP), (FP) also has the

---

sometimes describes his notion of exact verification as embodying a constraint of holistic relevance in the sense that for a state to exactly verify a proposition, it must be wholly relevant to the proposition, and so must not have any part that is irrelevant to the proposition (cf. e.g. (Fine, msa: p. 1)). This may invite the worry that some sort of untoward circularity is involved in using Fine's framework to describe and study relations of relevance. However, the appeal to a notion of relevance is confined solely to Fine's informal commentary on his theory, and not part of the theory itself. It might perhaps still be claimed that to the extent that exact verification imposes relevance constraints, an analysis of relevance within the truthmaker framework is in that sense not fully reductive; I would be content to concede that much.

[16] This is a slight exaggeration, since other treatments of negation are possible within truthmaker semantics that do not require a separate specification of falsifiers, appealing instead to modal connections on the states. These approaches to negation will not be considered here.

consequence that for any propositions P, Q, both P and Q are parts of $P \wedge Q$. Indeed, we might say that on (FP), being a part of a proposition is the same as being a conjunct of the proposition, for the condition that P is a part of Q turns out to be equivalent to the condition that $P \wedge Q = Q$ (cf. (Fine, 2015: p. 13f, 19)). In particular, then, on (FP), the conjuncts $P \vee \neg Q$ and $\neg P \vee Q$ of the hypothesis $(P \vee \neg Q) \wedge (\neg P \vee Q)$ in our example are classified as parts of that hypothesis. As a result, by interpreting the notion of part invoked in Gemes' criterion (GI) in terms of Fine's hyperintensional explication, we avoid the counter-example of the previous section.

4. Fine-tuning

I turn now to some difficulties for (GI) that do not arise from an inadequate conception of partial content and show how the criterion may be refined to avoid them.

4.1. **Parts of disjuncts.** In this section I argue that even under the Finean conception of parts of contents, (GI) overgenerates irrelevance, and propose a fix that employs a disjunctive counterpart of the notion of content parts. We begin by considering a slight variation on the case discussed in section 2, namely the following pair of evidence and hypothesis:

> **(E_3):** A came up heads.
> **(H_3):** (A came up heads and B came up heads) or (A came up tails and B came up tails).

Note that the hypothesis H_3, under the obvious formalization $(P \wedge Q) \vee (\neg P \wedge \neg Q)$, is logically equivalent to the previous hypothesis H_2, formalized as $(P \vee \neg Q) \wedge (\neg P \vee Q)$. As before, $E_3 = P$ is intuitively relevant to H_3, even though it does not lower or raise its probability. If we ask why E_3 appears relevant to H_3, the most natural answer goes roughly along the following lines: Firstly, E_3 guarantees the truth of one conjunct of the first disjunct of H_3, and thereby makes it more probable that this disjunct obtains. Secondly, E_3 rules out the truth of the first conjunct of the second disjunct of E_3, and thereby ensures that this disjunct does not obtain.

So again, the intuitive verdict of relevance may be seen to arise from a probabilistic effect of the evidence on propositions that are intimately related, though not identical, to the hypothesis. However, in this case, the propositions in question are *not parts* of the hypothesis on either Gemes' or Fine's account of partial content.[17] The disjuncts

---

[17] To see this, it suffices to note that none of these propositions – the disjuncts of H_3 and their conjuncts – are even logical consequences of H_3. On Gemes' view, content parts are by definition a special kind of logical consequence. On Fine's view, this is clear from the fact that content parts are conjuncts of the propositions they are part of.

of H_3 may instead be described as different *ways* for H_3 to hold, and their conjuncts accordingly as parts of ways for H_3 to hold.

Under this diagnosis, the present counter-example to (BI) is suggestive of a different revision of (BI) than that proposed by Gemes. Specifically, the case seems to suggest that if the evidence, or a part of it, is probabilistically relevant to a way for the hypothesis to hold, or perhaps even just to a part of a way for the hypothesis to hold, then this is sufficient for relevance.[18] Put in terms of irrelevance, the envisaged conditions read as follows:

(W): A piece of evidence E is irrelevant to a hypothesis H only if for every way H^* for H to hold, $\Pr(H^*|E) = \Pr(H^*)$.

(WP): A piece of evidence E is irrelevant to a hypothesis H only if for every part H' of some way H^* for H to hold, $\Pr(H'|E) = \Pr(H')$.

The notion of a way for a proposition to be true may then be defined as the disjunctive counterpart of the notion of partial content, so that H^* is a way for H to hold just in case $H^* \vee H = H$. In the framework of truthmaker semantics, this is in turn equivalent to the following definition, paralleling in an obvious way the definition of partial content:[19]

(FW)   *P* is a way for *Q* to hold iff (i) every verifier of *P* is a verifier of *Q*

(ii) every falsifier of *P* is part of a falsifier of *Q*

(iii) every falsifier of *Q* has a part that falsifies *P*

Admittedly, given the counterpart of (W) for parts

(P): A piece of evidence E is irrelevant to a hypothesis H only if for every part H^* for H to hold, $\Pr(H^*|E) = \Pr(H^*)$.

which is implied by (GI), we are not *forced* by the above example to accept either of (W) or (WP). The reason is that under the Finean interpretation of 'part', (P) already implies that E_3 is relevant to H_3, for the propositions $P \vee \neg Q$ and $\neg P \vee Q$ turn out to be parts of $H_3 = (P \wedge Q) \vee (\neg P \wedge \neg Q)$. But there is no reason to suppose that in general, evidence rendered relevant by (W) is also rendered relevant by (P) or (GI). The pertinent cases concern hypotheses that may be written as disjunctions $H_a \vee H_b$, where the evidence E is probabilistically relevant to each disjunct in such a way that its effects on the disjuncts

---

[18] There is no corresponding motivation to also take into consideration mere ways for the evidence to hold, or mere parts of such ways. Indeed, for any hypothesis H with $0 < \Pr(H) < 1$ and arbitrary P, $(H \vee \neg H) \wedge P = (H \wedge P) \vee (\neg H \wedge P)$ would otherwise turn out relevant to H on the strength of the probabilistic relevance of H to H. This would seem a bad result. Surely, amassing evidence of this sort by procuring arbitrary information P would amount to an objectionable waste of resources in an inquiry into H.

[19] Cf. (Fine, msa: p. 16); Fine says that *P exactly entails Q* when I say that *P is a way for Q to hold*.

cancel each other out. There is no reason to infer from this that either E or $H_a \vee H_b$ even *have* proper parts, let alone ones that are probabilistically relevant to each other. I conclude that we have strong reasons for accepting (W).[20]

On the basis of (W) and (P), we can now give an argument for the stronger claim that (WP). For given (W), it is plausible also to accept the strengthening on which the mere relevance (probabilistic or otherwise) of the evidence E to some way H^* for the hypothesis H to hold is sufficient for E's relevance to H. But then suppose that E is probabilistically relevant to part H′ of the way H^* for H to hold. Then by (P), E is relevant to H^*. By the strengthening of (W), it follows that E is relevant to H, just as required for (WP).[21]

4.2. **Helpful parts.** I shall now give an argument that (WP), and even (P), overgenerates relevance, and propose a modification that avoids the problem. Crucial to the argument and the modification is a distinction between what I will call *helpful* and *unhelpful* parts of a way for a proposition to hold. It will help to have a short term for parts of ways for a proposition to hold, so alluding to their status as *c*onjuncts of *d*isjuncts, I will call them cd-parts.

Consider some proposition P and assume that state s is not a verifier of P, but that s is a proper part of some verifier t of P. Then s, we might say, goes some way towards making P true, though not the whole way. However, note that it may still be the case that, as it were, the truth of P is as far away given the state s as it is without s. For we may ask what states can be fused with s so as to yield a verifier of P. And it may be that the only states satisfying this description are themselves already verifiers of P. If so, then in terms of what is still needed to make P true, the state s does not bring us any

---

[20] Note that (W) implies that whenever $0 < \Pr(P) < 1$, P is relevant to $P \vee \neg P$, since P is then probabilistically relevant to P. This is in marked contrast to (BI), and to (GI) on Gemes' account of content parts, on which nothing can be relevant to a logical truth. Since $\Pr(P \vee \neg P)$ is always 1, $P \vee \neg P$ makes for a somewhat peculiar choice of a hypothesis to investigate, so it is not obvious what significance to attach to our result. However, if we wish to allow for rational inquiry into a hypothesis that is a logical truth like $P \vee \neg P$, then the result seems very plausible to me. For P then *is* evidence that bears on $P \vee \neg P$ in a way in which it does not bear on arbitrary $Q \vee \neg Q$, and it seems to me a feature, not a bug, of the present proposal, that it enables us to capture this fact.

[21] It is plausible that (P) may be strengthened in the analogous way, so that consequently E's probabilistic relevance to a way for H′ to hold, where H′ is part of way for H to hold, is also sufficient for E's relevance to H. Fortunately, this is already implied by (WP). For in this case, H may be written $((P \vee Q) \wedge R) \vee S$, where E is probabilistically relevant to P. We can then show that $P \wedge R$ is a way for H to hold, and thus P part of a way for H to hold, using that $(P \vee Q) \wedge R = (P \wedge R) \vee (Q \wedge R)$, and thus H $= ((P \wedge R) \vee (Q \wedge R)) \vee S = (P \wedge R) \vee ((Q \wedge R) \vee S)$. The identities used here are implicit in the soundness results of (Fine, 2015: sections 6 and 9).

closer to the truth of *P*, even though it goes some of the way towards making *P* true. Now for evidence to be relevant to a hypothesis, I want to suggest, it (or one of its parts) has to be probabilistically relevant to not just any cd-part of the hypothesis, but to one that brings us closer to the hypothesis.

Consider the following example, concerning again a throw of two dice A and B.

(E$_4$): B came up 1.
(H$_4$): A came up even or (A came up even and B came up odd).

Let *P* be the proposition that A came up even, and *Q* the proposition that B came up odd, so H$_4$ = $P \lor (P \land Q)$. Note that the verifiers and falsifiers of *Q* stand to truth and falsity of H$_4$ in the way just described. For consider any state, such as B having come up 1, or 3, or 5, that verifies *Q*. For any such state *s*, we may now ask what we can add to that state so that we obtain a verifier of H$_4$. We then see that the only states of this sort are themselves already verifiers of H$_4$, namely the states of A having come up 2, or 4, or 6. Likewise consider any state, such as B having come up 2, or 4, or 6, that falsifies *Q*. For any such state *s*, we ask what we can add to it so as to obtain a falsifier of H$_4$. We then see that the only states of this sort are themselves already falsifiers of H$_4$, namely the states of A having come up 1, or 3, or 5.

Should we consider E$_4$ relevant to H$_4$? E$_4$ is not probabilistically relevant to H$_4$, which can be seen from the fact that H$_4$ is logically equivalent to *P*. But since *Q* is a cd-part of H$_4$ and E$_4$ is probabilistically relevant to *Q*, E$_4$ is classified as relevant to H$_4$ by (WP). Moreover, H$_4$ = $P \lor (P \land Q) = P \land (P \lor Q)$, hence $P \lor Q$ is part of H$_4$. Since E$_4$ is probabilistically relevant to $P \lor Q$, E$_4$ is already classified as relevant to H$_4$ by (P). It seems to me that this is the wrong result. If you were called upon to investigate the hypothesis H$_4$, it would appear rational for you to discard the information E$_4$ as irrelevant, and you might rightly be blamed for wasting time if you were to spend it on procuring the information E$_4$. It is natural to take the reason for this to be that the information E$_4$ does not get you any closer to an answer to H$_4$. The above observations concerning the verifiers and falsifiers of E$_4$ and H$_4$ give a precise sense in which is true.

Since the hypothesis H$_4$ has a somewhat contrived and unnatural logical structure, intuitions about the example are perhaps less firm than we should like them to be to motivate replacing (WP) by (HELP). So let me try to marshal some additional support for this move. Suppose you are interested in a hypothesis H, but it is difficult to obtain any evidence that bears probabilistically on the hypothesis taken as a whole. So you move to considering parts of ways for the hypothesis to hold, to see if data bearing on the probability of these might be more easily obtained. Then (HELP) can be seen as

amounting to the following, very reasonable injunction: Make sure, in selecting a cd-part P of H to collect data on, that any way for P to be true brings us closer to the truth of H, and that any way for P to be false brings us closer to the falsity of H.

Let us state the proposed modification of (WP) more explicitly. First, we define what it is for a proposition P to bring us closer to another Q, or as I shall say, for P to *help* Q.

(H) P helps Q iff for every $s \in P^+ \setminus Q^+$, there is a $t \notin Q^+$ with $s \sqcup t \in Q^+$, and

for every $s \in P^- \setminus Q^-$, there is a $t \notin Q^-$ such that $s \sqcup t \in Q^-$

It can be shown that if P helps Q according to this definition, this ensures that P is part of way for Q to hold.[22] We may therefore replace (WP) above with this weaker alternative:

> **(HELP):** A piece of evidence E is irrelevant to a hypothesis H only if for every helper H′ of H, $\Pr(\text{H}'|\text{E}) = \Pr(\text{H}')$.

Here, then, is my proposal for a criterion of irrelevance that fits the role intended for the notion of irrelevance in separating what may be discarded and what should be valued in a given context of inquiry:

> **(IRRE):** A piece of evidence E is irrelevant to a hypothesis H iff for every part E′ of E and every helper H′ of H, $\Pr(\text{H}'|\text{E}') = \Pr(\text{H}')$.

5. Conclusion

Gemes has argued convincingly that there is an important distinction between relevant and irrelevant evidence that is not adequately captured by the usual Bayesian criterion. However, his own proposal, while a significant improvement, still has unacceptable consequences. To do better, I have argued, we need to accept that the notion of relevance is *hyperintensional*; it is sensitive to differences in content that may obtain even between logically equivalent propositions.[23] I have then utilized Fine's framework of truthmaker

---

[22] We can construct a proposition R of which P is part and which is a way for Q to hold as follows. Let the set of verifiers of R be the set of verifiers of Q that have a part which verifies P. Let the set of falsifiers of R be $(P^- \cup Q^-)^\circ$. It is then straightforward to show that R is a proposition, and that it relates to P and Q in the desired way.

[23] In this respect, the present paper would seem to follow something of a trend. Hyperintensional accounts have in recent years been proposed for many philosophically central concepts, such as essence, ground, conditionals, subject matter, and, closest to our present concerns, confirmation – cf. here esp. Yablo (2015). It is striking that in many cases, considerations of relevance play an important role in motivating the claim to hyperintensionality. It would be very interesting to explore the connections between these debates and the arguments I have here advanced in detail. A particularly tight connection may obtain to *ground*, for which Fine has offered a semantics within the same truthmaker framework we have employed here (cf. Fine (2012a,b)). Indeed, our notion of a cd-part coincides with Fine's

semantics to formulate a hyperintensional criterion of irrelevance that avoids the difficulties that befell the intensional account of Gemes. In view of the advantages the resulting notion of irrelevance enjoys over its Bayesian and Gemesian rivals, it would be very interesting to develop the theory of this notion in detail, determining its formal properties and its relation to other, related notions definable within the truthmaker framework. This task, however, I have to leave to future work.

REFERENCES

Cohen, L. J. (1994). Some steps towards a general theory of relevance. *Synthese 101*, 171–185.

Fine, K. (2012a). Guide to ground. In B. Schnieder and F. Correia (Eds.), *Metaphysical Grounding*, pp. 37–80. Cambridge: Cambridge University Press.

Fine, K. (2012b). The pure logic of ground. *The Review of Symbolic Logic 5(1)*, 1–25.

Fine, K. (2013). A note on partial content. *Analysis 73(3)*, 413–419.

Fine, K. (2015). Angellic content. *Journal of Philosophical Logic Online First*, 1–28.

Fine, K. (msa). A theory of truth-conditional content I. Unpublished manuscript.

Fine, K. (msb). A theory of truth-conditional content II. Unpublished manuscript.

Gemes, K. (1994). A new theory of content I: Basic content. *Journal of Philosophical Logic 23*, 596–620.

Gemes, K. (1997). A new theory of content II: Model theory and some alternatives. *Journal of Philosophical Logic 26*, 449–476.

Gemes, K. (2007). Irrelevance: Strengthening the Bayesian criteria. *Synthese 157(2)*, 161–166.

Keynes, J. M. (1929). *A Treatise on Probability*. Macmillan.

Yablo, S. (2015). *Aboutness*. Princeton University Press.

---

notion of a weak partial ground, and our notion of a way for a proposition to hold coincides with Fine's notion of a weak full ground.