

Implication, Evidence, and Subject Matter

Stephen Yablo

Received: date / Accepted: date

Abstract Insert your abstract here. Include keywords, PACS and mathematical subject classification numbers as needed.

Keywords First keyword · Second keyword · More

1 Closure

If Alma knows that p , and p implies q , must Alma know that q ? Of course not. She may not even understand q . She may understand it, but not believe it. She may believe it for the wrong reasons, inferring q from a premise that she doesn't know (rather than p , which she does). She may be reasoning from the right premise, but in a bad way.

But, while explanations can be imagined of Alma's ignorance, that an explanation seems called for is interesting in itself. Alma *could* have known that q , it seems, if she had played her cards right:

(KC) If A knows that p , then A knows, or is in a position to know, that q .

Note the disjunction: she knows, or is in a position to know. The second disjunct adverts to the possibility just mentioned, of Alma parlaying her knowledge of p into knowledge that q . She doesn't need to do any parlaying, if she already knows that q ; the first disjunct is in recognition of the possibility that Alma is already there.

F. Author
first
Tel.: +123-45-678910
Fax: +123-45-678910
E-mail: fauthor@example.com

S. Author
second address

The closure principle seems right when we consider it in the abstract, but our confidence may waver when we look at particular instances. The second claim *seems* in each case to be less knowable than the first:

(Moore) I have a hand. There is an external world.

(Nozick) I am sitting by the fire. I am not a bodiless BIV.

(Dretske) That is a zebra. It is not a cleverly disguised mule.

(Kripke) I turned the stove off. Evidence that I didn't is thus misleading.

(Cohen) It is 3am (says my watch). My watch is accurate iff it reads 3am.

(Vogel) I am teaching logic next year. So, I won't die in the meantime.¹

Never mind whether the examples are bogus or genuine. What there does seem to genuinely be is a phenomenon of *apparent* closure violations. Either we feel the pull ourselves, or can tell when it is apt to be felt by others. There is something here that tempts us to *think* that closure is violated, whether we give in to the temptation or not.

Kripke worked in *Naming and Necessity* with a notion of appearances, or intuitions, of possibility—IPOs, let's call them. He asked how they arise, on the theory that an IPO loses credibility if the explanation doesn't require it to be true. Let us speak in a similar spirit of intuitions, or intimations, of openness—IONs, for short. Asking how they arise may, as with Kripke, help us decide what to make of them.

To be sure, a lot of epistemologists already have decided. They see IONs not as data to be reckoned with, but anomalies to be explained away. I do not pretend to know all the reasons for this attitude. But one important reason is this: difficult as closure may be to accept in some cases, there seems to be no good way of denying it.

What would be a good way of denying closure? A principle with this kind of pedigree cannot just be thrown under the bus. A good way of denying it would tell us what was right in the principle—call that the “defensible core”—and explain how the remainder can be done without. This proves difficult.

One problem is that no one has any idea of what the defensible core is supposed to be. A second is that it would have to be extremely weak, since the full principle can be regained from modest assumptions. Third, for closure to be restricted *at all* means that deduction is unreliable, taking us sometimes from known premises to unknown conclusions. Not having been on guard against this possibility, conclusions we don't know are thoroughly interwoven in with ones that we do. There is no sorting them out now; so the edifice of knowledge will have to be reconstructed from scratch. No one can take this seriously.

If this is where the debate stands, we see why it is found frustrating. One side—the losing side, at present—lays great weight on a few intuitively vivid anomalies. The other side does not deny the anomalies! They just do not see them as anything like decisive, in the absence of a workable alternative. This is what make Sharon and Spectre's evidential approach so interesting. The logic of evidence holds promise of stirring things up, by showing how it could come about that one comes to know that p without coming to know that q .

¹ [?], [?], [?], [?]

2 Support

Closure principles can be formulated for lots of notions. Along with closure for knowledge—closure_k—there is closure for warrant, certainty, and so forth. S&S want us to look at closure_s: closure with respect to *evidential support*.

(SC) If e supports p , then it supports any consequence q of p .

The fate of knowledge-closure is tied up, in their view, with that of support-closure. This is for two reasons a beautiful idea. What Alma knows depends, obviously, on the evidence she possesses. If her evidence “gives out” between p and q , her knowledge might be expected to give out too. Well, but does her evidence give out in this way? Of course, that question could be as recalcitrant as the corresponding question about knowledge, in which case we will be no further ahead. But, and this is the second reason it’s a terrific idea, the theory of evidential support is, thanks to Hempel and others, far more advanced on this score than the theory of knowledge.

Hempel singled (SC) out for attention already in the 1940s. *Special Consequence*, as he called it, was pitted in his work against *Converse Consequence*, which corresponds more or less to the hypothetico-deductive model of confirmation.² Goodman in *Fact, Fiction, and Forecast*³ subjected *Special Consequence* to further (withering) scrutiny. Carnap, when he attempted to read the laws of evidence off the laws of probability,⁴ noticed that probability theory did not support any such principle as (SC). It makes indeed for *counterexamples* to support-closure. If evidential support is to be understood as probability-raising, then it is very definitely not closed under entailment. This sets up the main question of S&S’s paper:

if the openness of evidence can be established..., and some kind of dependence of empirical knowledge on evidence is unavoidable..., how can knowledge be closed? (2)

Their basic argument is sketched in section ??; it features a scenario (the “breakout scenario”) whereby failures of support-closure appear to make for failures of knowledge-closure.⁵ Overgeneration worries are raised in section ??, and connected with the issue of whether closure should be expected to fail in section ?. Section ??. Section ?? argues that q , to play its assigned role in the breakout scenario, must already be “almost” known; it resembles a Gettier or lottery proposition in being a truth believed on the basis of evidence that, although fully justifying, not the right type for knowledge. Sections ?? and ?? question whether “knowledge-conducive evidence” is really lacking in these cases, and

² [?,?]

³ [?]

⁴ [?]

⁵ S&S devise multiple interlocking objections to support-closure and shoot down many imagined responses. One gets the feeling that (SC) is the territory we are primarily fighting over. That is territory long ago surrendered, though, I would have thought. The focus should be on whether evidential openness carries over to knowledge.

more generally whether the observed pattern of IONs is what one would expect if the problem was knowledge-conducive evidence. Section ?? introduces immanent closure: if Alma knows that p , then she knows, not perhaps all its consequences, but its parts. Section ?? defines the notion of part: q is part of p if p includes it both truth-conditionally— p entails q —and in what it's about— q 's subject matter is included in that of p . It is argued in section ?? that closure should be expected to sometimes fail if q addresses matters on which p is silent, even if q is implied by p . The context-sensitivity of knowledge is linked in section ?? to the discourse-sensitivity of sentential subject matter. Section ?? considers the discourse context where q is inferred from p . An appendix tracks some of these ideas back to Hempel, and uses them to defend a version of special consequence, or support-closure.

3 The breakout scenario

“[I]f the openness of evidence can be established..., and some kind of dependence of empirical knowledge on evidence is unavoidable,...how can knowledge be closed?” (2). Well, the dependence might be of a kind that *allows* it to be closed. Knowledge might, for instance, require high evidential probability, which extends from p to its consequences, rather than evidential support, which does not. S&S lay out a scenario, however, which appears to establish the desired connection. I call it the *breakout scenario*, for this reason. Confirmation has been studied mainly by *formal* epistemologists and *scientific* epistemologists. The scenario has them breaking out of their tech-y shell to begin making pronouncements about *knowledge*, until now the province of regular epistemologists.

BREAKOUT SCENARIO Alma starts out not knowing p , and not knowing its consequence q either. She encounters evidence e that supports p but not q . The evidence is good enough that she winds up knowing that p . Closure requires her also to know q . But e is irrelevant to q , or even negatively relative to it. Alma cannot come to know q on the basis of evidence pointing away from q ! She gains knowledge that p but not that q , which is contrary to Closure.⁶

The principle S&S are relying on in the second last sentence is

Evidence dependence (ED): If A does not know at time t_0 that q , and the evidence she acquires between t_0 and t_1 counts against q , then A does not know q at t_1 either.

To see the problem with this principle, recall the Bayesian problem of old evidence. To count in favor of h , we are told, e must boost its probability. This

⁶ “A subject S comes to gain evidence e for p (evidence which does not a priori entail p) on the basis of which S comes to know that p is true. Since for all p there will be propositions q which a priori follow from p but are not supported by e , if S did not know q before the evidence came in, S does not (given ED) know it after” (p. 24).

will be difficult, if we were already certain of e , for $\text{pr}(h|e) = \text{pr}(h)$ when $\text{pr}(e) = 1$. Can e 's value as evidence really depend on our not having acquired it too early?

The principle used in the breakout scenario is analogous. To count as evidence for h , anyway the kind to confer knowledge, e must increase Alma's confidence in h . This will be difficult if Alma was already highly confident of h . If Alma's attitude to h was based on evidence that is undermined when she encounters e , then h 's probability may well drop, even if e would ordinarily suffice for knowledge. It makes no sense that e 's claim to be regarded as good, knowledge-conferring evidence depends on Alma not having earlier encountered highly seductive bad evidence.

I can illustrate with a Gettier proposition, though the problem is more general. Let the hypothesis be *Jones own a Ford or Brown is in Barcelona*. Alma accepts it, initially, because of misplaced confidence in Jones owning a Ford. Brown writes, saying, "News of your mistake about Jones has reached me in Barcelona." Alma didn't know h to begin with, but she does now, since her confidence is well founded. Ah, but the new confidence may be slightly less confidence-inspiring than her ill-founded old confidence in Jones's owning a Ford. (ED) will in that case object, but it shouldn't. (ED) says in effect that one can never live down early misadventures with the wrong disjunct. I may never recover, for instance, from being convinced by my parents that the car belonged to Jones, no matter how many telegrams eventually arrive Brown.⁷ Meanwhile others can know based on the exact same telegrams.

This may seem far from S&S's concerns. But it fits some of their cases fairly well. Observation reveals that a book is red, but not, supposedly, that it is red if it looks red. Rewrite the conditional as a disjunction: *Either it does not look red, or it is red*. Initially, in advance of observation, we are drawn to the first disjunct more than the second; the book is just as likely to look green, or blue, etc. Observation shifts our confidence to the second disjunct, the true one. The book's appearance gets us believing a truth for the right

⁷ The sins of the fathers are visited upon the sons.

reason, the reason one should be relying on if one wants to know.⁸ Having raised this worry, let me put it aside for now.

4 Unintended instances

Alma's evidence e is meant to give her knowledge of p , while leaving her still ignorant of whether q . This is a tricky business. If q might, for all Alma knows, be false, then it is natural to suppose that p too might, for all Alma knows, be false. "For all she knows" is a kind of relative possibility. If k is what Alma knows, we're talking about the possibility of $k\mathcal{E}^{\mathcal{E}}\neg q$ and $k\mathcal{E}^{\mathcal{E}}\neg p$. A world witnessing the first possibility ipso facto witnesses the second.

Here is the kind of case that S&S like: Alma watches Bolt win the 100 meter dash on TV, reads about it in the newspaper, etc. The evidence she gets from these sources does not address the issue of what will become of Bolt's refrigerated blood sample. Testing could reveal that Bolt had been using a banned substance. He would in that case be disqualified, with the result that he will not technically have competed or hence won the gold. Alma does know that Bolt is the Olympic champion, it seems. Does she know that blood tests won't be devised on the basis of which Bolt is disqualified? Surely not; she has no evidence on that issue. The case is somewhat believable.

Their apparatus allows, however, for a different case: Albert learns from TV and newspapers that Bolt has won the gold and Blake the silver. The evidence he gets from these sources is all about Bolt. It does not address the issue of whether Blake won the silver. That is OK, though; knowledge that p does not require one to tick off all its necessary conditions q . Just as Alma needn't know about future blood tests, Albert needn't know about Blake. This is hard to make sense of. For Blake to win the silver is not one more necessary condition. It's part of what it *is*, we want to say, for Bolt to win the gold and Blake the silver.

⁸ Another way to see the problem with (ED). Evidence e that refutes b surely also refutes a hypothesis c that strictly entails b . (ED) cannot allow this when $c = e\&b$; e as a consequence of $b\&e$ cannot lower its probability. It seems like double-counting for e 's recurrence in a hypothesis to be what spares the hypothesis from refutation by e . The principle that suggests itself, letting $p\text{-}q$ be what remains when q is extricated from p , is this.

(1) e is evidence for (against) p iff e makes $p\text{-}e$ likelier (less likely).

Call that the remainder principle. Putting $e\mathcal{E}^{\mathcal{E}}\neg h$ for p ,

(2) e is evidence for (against) $e\&\neg h$ iff e raises (lowers) the probability of $(e\&\neg h)\text{-}e$.

Putting $\neg h$ for $(e\&\neg h)\text{-}e$,

(3) e is evidence for (against) $e\&\neg h$ iff e raises (lowers) the probability of $\neg h$.

Evidence against $e\&\neg h$ is evidence for its negation $e\supset h$, so

(4) e is evidence for (against) $e\supset h$ iff e raises (lowers) the probability of h .

If and when the remainder principle holds, e 's evidential relation to $e\supset h$ is the same as its relation to h .

I draw the following moral. To know that p , it is not enough to know *some* of what it says; one must know all of it. This applies when q is a conjunct of p , but not in the original case, because *Bolt took the gold* only implies—it does not in part *say*—that Bolt won’t be disqualified.⁹

Knowing that Bolt won, but not about his blood sample, certainly *sounds* less ridiculous than knowing he and Blake came in first and second, but not that Blake came in second. But there might be various reasons for this. Maybe it sounds bad to *claim* knowledge that Bolt will not be disqualified, even if the knowledge is there. Maybe one cannot know on the *basis* of his winning the gold that he will never be disqualified.¹⁰ (Explanations along these lines have been suggested for many, if not all, of the best known prima facie counterexamples.) Or, the counterexample may be genuine. It is hard to say at this point.

Sharon and Spectre might say we are getting ahead of ourselves, dismissing the appearances before we have even tried to make sense of them. It does seem that Alma can learn Bolt won the gold from evidence silent on the issue of subsequent disqualification. And it does not seem that Albert can learn that Bolt won gold and Blake silver from evidence leaving him ignorant of the second conjunct. If this is right, then whatever it is that enables Alma to learn only p in the first case, is not operative in the case of Albert. This doesn’t yet tell us what the explainer is, of course. But it suggests a way to find out. Closure fails because of a factor that is present in Alma-type instances of the scenario, and absent in Albert-type instances. If a factor like that comes to light, we will want to investigate it further, to see how well its operations line up with the intuitive data on IONs generally.

5 Explanatory desiderata

Sharon and Spectre maintain that closure violations are unsurprising—they “make sense”—when we bear in mind that evidence for p does not always carry through to q . What does it mean for them to “make sense” from an evidential perspective? It might mean just that we would expect closure violations to occur, viewing the matter from that perspective. But everything makes sense from *some* perspective. Who is to say that perspective is the right one to take? One has to look at the mechanism that it postulates. Is the mechanism

⁹ Various intuitive tests confirm this. Alma in asserting that Bolt has won the gold cannot be said to have asserted inter alia that he won’t be disqualified. Albert, in asserting that Bolt won the gold and Blake the silver, *has* asserted inter alia that Blake won the silver. If Blake turns out to have come in fourth, still Albert was not wholly wrong; his statement is partly true by virtue of what it says about Bolt. But what if it is Bolt who turns out to have come in fourth? Is Alma partly vindicated? No, she was entirely wrong. The fact that Bolt was not later disqualified does not confer partial truth on *He won the gold*.

¹⁰ Compare *Bolt won the gold* and *He won the gold and will never be disqualified*. These are a priori equivalent, so they ought to be equiknowable. The scenario does not bear this out. If to learn that Bolt won the gold and Blake the silver, Bina must know that Blake won the silver, then, it seems to me, to learn that Bolt won the gold and will not be disqualified, Alma must know that he will not be disqualified.

believable? Are the expectations it generates as to when closure should fail actually borne out?

Prima facie counterexamples to (KC) have been around for a long time. They have come to be seen, as remarked in section ??, not as data to be reckoned with, but anomalies to be explained away. Some take this attitude in the belief that most IONs *have* been explained away.¹¹ Some don't want to jeopardize knowledge by deduction. But there is also a Kuhnian reason for not jumping ship. We have nowhere to jump *to*.

This may seem over-dramatic. Certainly ways have been found to *contain* closure. But they don't contain it in the right sort of way. All existing technologies wind up either strangling closure in the cradle, or leaving too much of it in place. Nozick's approach is the most familiar, so let's focus on that. Nozick thinks closure ought to fail when q is a skeptic-baiting consequence (*I am not a brain in a vat*) of some evident truth (*I am sitting by the fire*). He offers a conception of knowledge that seems at first to deliver this result:

(CF) S knows that p just if S would not have believed it, had it been false.

I would not have believed I was sitting, had I been standing, or in a crouch. If I were a brain in a vat, however, my beliefs would be in relevant respects unchanged. If I were not going to be teaching logic next year, I would not now expect to be teaching it. Suppose, though, that I was going to be hit by lightning. Would I have seen it coming? Of course not. Acts of God are not the kind of thing one realizes in advance.

These might be regarded as acceptable violations of closure. Kripke points out, however, that (CF) makes as well for *intolerable, egregious* violations of closure ([?]). To know that there is a red barn in the field is not necessarily to know there's a barn there—for the closest alternative to a red barn might be a green one, while the closest alternative to a barn was a perfect red duplicate.

Hawthorne notes that (CF) makes as well for egregious *non*-violations; it is not as restrictive as its advocates intended. Granted that I do not know, on Nozick's view, that I'm not a brain in a vat, I do know I am not a brain in a vat in a dark room, just by virtue of knowing that I am not in a dark room.¹²

If Nozick had been able to provide a unified explanation of all IONs, that would have given closure fans something to worry about. They would have had to come up with a unified account of why closure only seems to fail in these cases. As it is, they are free to explain the examples piecemeal, each in its own way. The same dialectic obtains here. S&S are suggesting a new closure-containing technology. This puts pressure on closure-fans, *if* S&S can

¹¹ Perhaps as failures of warrant-transmission ([?], [?], [?], [?]).

¹² [?] makes this point against Dretske. Both sorts of egregiousness have same source, that if (CF) grants me knowledge of r , it grants me knowledge as well of any $r\mathcal{E}s$ such that r fails in nearer-by worlds than s . In Kripke's example, r is *It is red* and $r\mathcal{E}s$ is *It's a red barn*. In Hawthorne's, r and s are *I am not in a dark room* and *I am not a brain in a vat in a dark room*.

pull under one umbrella phenomena that closure fans are forced to treat separately. Otherwise, maybe not. One would like to know, not only how evidential considerations could lead to closure violations, but how well they fit the observed *pattern* of IONs.

6 “Foreknowledge”

The breakout scenario is built on a failure of support closure. The evidence e that carries p over the knowledge threshold is irrelevant, or negatively relevant, to q . The situation rational-confidence-wise is this:

1. e raises A 's confidence in p
2. e doesn't raise A 's confidence in q ¹³

The scenario also appears to involve, taking p 's truth as given, that

3. p is initially unknown, because $\text{pr}(p)$ is too low for knowledge
4. q is initially unknown, presumably because $\text{pr}(q)$ is too low for knowledge¹⁴
5. p is ultimately known, because $\text{pr}(p|e)$ is high enough for knowledge
6. q remains unknown, presumably because $\text{pr}(q|e)$ is too low

The killer combination intuitively speaking is (2) with (4). That q , already in trouble by (4), is made, if anything, *less* probable by our newly acquired evidence feels like the final insult. A disconfirmed hypothesis of which we were already ignorant would seem to have not much going for it.

There is, we noted, another way to view the matter. q 's decline may just reflect how probable it was to begin with.¹⁵ It had, in the circumstances, nowhere to go but down. Take the case that S&S lead with and use as a model for the rest: q is $e \supset p$. If $\text{pr}(p|e)$ is high enough for knowledge, as the scenario requires, then $\text{pr}(e \supset p)$ is high enough too; $\text{pr}(e \supset p)$ is never less than the conditional probability and usually more. Using “knowledgeably confident” for the confidence level required for knowledge,

Learning that e does not give Alma knowledge of p , unless she was knowledgeably confident of $e \supset p$ beforehand.

Indeed any consequence of q of p that is not supported by e has this property:

Learning that e cannot give Alma knowledge of p , unless she was knowledgeably confident of q beforehand.¹⁶

¹³ Update is Bayesian throughout: p 's new probability, or learning that e , is the conditional probability beforehand of p on e .

¹⁴ No other reason is given, at any rate, why A should fail to know, initially, that she will not inherit a fortune, or that Bolt will never be disqualified.

¹⁵ “Probability” is always rational probability, the degree of confidence we're entitled to.

¹⁶ If q were not sufficiently likely at the beginning, p would not be sufficiently likely at the end. $\text{pr}(q|e) \leq \text{pr}(q)$ since e fails to confirm q ; $\text{pr}(p|e) \leq \text{pr}(q|e)$ since p implies q . $\text{pr}(p|e)$ is p 's final probability, though. If $\text{pr}_{old}(q)$ is too small, then $\text{pr}_{new}(p)$ is too, for it's smaller.

But then, our description above was not quite right. What really happens is this:

1. e raises A 's confidence in p
2. e does not raise A 's confidence in q
3. A is not at first knowledgeably confident of p ; that's why A doesn't know
- 4'. A fails to know q , at first, although she *is* knowledgeably confident of q
5. A winds up knowledgeably confident of p , enabling her to know p
- 6'. A winds up not knowing q , despite being more confident of it than of p ¹⁷

From this we see that the breakout scenario does not go through if a truth that Alma is rightly confident of is a truth she knows. Alma would know q throughout, if justified true belief sufficed for knowledge.

Of course, we know from Gettier, lottery, and (on some views) anti-skeptical propositions, that truth and high rational confidence do *not* suffice. But it pulls the sting, a bit, from Alma's ignorance of q , that it does not stem any lack of rational confidence or justification.¹⁸

7 Lack of evidence

Something prevents q , a truth of which Alma is properly confident, from being known by Alma, both as the scenario begins and as it ends. What is it? S&S say, "the reason we tend to deny the status of knowledge to the conclusions of such inferences is that they lack evidential support" (p. 3). Also that, "Lacking evidential support, empirical beliefs of this sort do not qualify as knowledge" (p. 6). Alma fails to know early and late—as the scenario begins and as it ends—so it must be that she lacks evidential support early and late.¹⁹

(1) Seeing a zebra-looking animal in the pen, although providing me with evidence that there is a zebra in the pen, does not provide any evidence that the animal is not a disguised mule (p. 5)

(2) Memory of having parked one's car in the driveway ten minutes ago evidentially supports the belief that one's car is in the driveway. It

¹⁷ Which she does know. $\text{pr}(q) > \text{pr}(p)$ on account of q being implied by p .

¹⁸ "It might be thought that since the probability of q cannot be lower than that of p , if p is known q must be known, or at least knowable, as well. But as the case of lotteries shows high probability conditional on the total evidence does not guarantee knowledge. Our argument concerns what one has evidence for, i.e. relative to any state what does one have evidence for given all of one's evidence, and not on the probability of propositions on one's total evidence or the degree of rational credence (which is influenced by initial credence assignments)" (section 4.2 (p. 60))

¹⁹ "One does not know at the outset that one's watch is accurate, that the car has not been stolen, or that the animal in the pen is not a disguised mule (we examine views to the contrary in the following section). The evidence one gains— by looking at the watch, recalling where the car was parked, or seeing a zebra-looking animal—counts against the truth of these propositions. Since counter-evidence cannot be the basis on which knowledge is gained, one does not know these propositions—although one can deduce them from what one knows" (p. 26).

provides no evidential support for the entailed belief that one's car has not been stolen in the last ten minutes (p. 6)

(3) Seeing what appear to be one's hands is evidence that one has hands; it is not evidence that one is not a bodiless brain in a vat (p. 6)

(4) Having proper evidence that p is true can allow one to know p , but not that evidence against p is misleading (p. 6).

(5) it is doubtful that by looking at one's watch, one can know that it is accurate. After all, the evidence one has gained counts against this conclusion (p. 10)

What we know about evidential support, at this point, is that it is lacking in examples like (1)-(5), and lacking when e makes q less probable. Even this, however, creates certain interpretive puzzles. Let me illustrate with (4) and (5). I don't doubt that they are IONs. The problem is that evidential support can be understood so that Alma *has* it in these cases.

Take first the Kripke dogmatism paradox, number (4) on our list. Evidence that p is true is said to provide no evidential support for $q = \textit{Evidence against } p \textit{ is misleading}$, that is, *Evidence against } p \textit{ is evidence against a truth}*. Hypotheses of this form are not intrinsically unknowable, I assume. I myself know that whatever evidence may turn up against my brother being named 'Paul' is evidence against a truth. How such knowledge is acquired is the question. I can imagine two approaches, one focused on the evidence against p , the other focused on p . On the first approach, we bracket the issue of whether p is true, and try to guess the *kind* of evidence there is liable to be against it. The presumed evidence is then examined for signs of misleadingness. Rumors cannot be trusted, the moon looks bigger on the horizon, water is not as cold as it feels, and so forth.

But, we typically have no idea what sort of evidence there may be against p . Even if that information were available, the approach seems strangely roundabout. The quick and easy way to find out whether anti- p evidence is evidence against a truth is to seek evidence directly on the question of whether p is true. S&S suggest that the quick and easy approach cannot work. First-order evidence for p is not evidence for the misleadingness of evidence against p . But they give no real argument for this claim, and it seems at least debatable. Suppose we were wondering whether evidence against p was evidence against something funny. An analogue of the first approach would be to look, say, at the newspaper that gave us this evidence, for signs of bias in favor of funny falsehoods. That doesn't sound very promising. To determine if evidence against p is evidence against something funny, we should ask directly whether p is funny. Why should a similar strategy not sometimes work for truth? Sharon and Spectre should find this plausible, given their acceptance of

Evidential Equivalence (E-EQ)

If e supports p , it supports any q that is a priori equivalent to p .²⁰

Evidence against p is evidence against a truth is just such a q . It implies that p , and is implied by p on the assumption that counterevidence always exists.

Number (5) on our list relates to the easy knowledge problem. A watch reading 3am is evidence for *It is now 3am*, but not, we are told, for *If the watch says it is 3am, then it is*. How better to establish that *Either my watch doesn't say so, or it is 3am* ($\neg e \vee p$), than to look at the watch? Seeing that it reads 3am tells me that the first disjunct ($\neg e$) is false, which means that q stands or falls with its other disjunct p : *It is now 3am*. By hypothesis, *The watch reads 3am* counts in favor of this disjunct. How can evidence for a hypothesis p that is known to agree in truth-value with q fail to be evidence for q ?

I am not questioning here that e makes q ($= e \rightarrow p$) unlikelier, or that e to confer knowledge should boost something's probability. What I am questioning is whether the something has to be q itself, as opposed to the disjunct p on which q 's truth-value turns out to depend? To bring some numbers into the picture, let the watch be digital. I am to begin with completely unaware of the time; none of the 1440 (24×60) is likelier than any other. My initial confidence in a reading other than 3am is $1439/1440 = 99.93\%$. That goes to zero when I look at the watch, and (given its one in a thousand failure rate) I wind up 99.9% confident that the time is 3am. That I had more confidence in $\neg e$ than I later develop in p seems irrelevant, since that original confidence was misplaced. I agree that (4) and (5) are IONs! We are asking how their ION-hood comes about. The problem does not seem, so far, to be a lack of knowledge-conducive evidence.²¹

8 Structural issues

Closure fans have explanations to offer of a good many IONs. Even if q is known, there are lots of reasons why it might seem more precarious than p . Perhaps

²⁰ "Although rejecting (E-EQ) provides a quick way out of the paradox of the ravens, the plausibility of the equivalence of evidence advises against this strategy" (p. 32).

²¹ *The watch reads 3am \rightarrow It is now 3am* may sound to us like it asserts a reliable connection. But this is plausibly for Jacksonian reasons ([?]). What *If e, then p* says, according to Jackson, is that $e \rightarrow p$. It is not assertible, however, unless $e \rightarrow p$ is robustly probable, that is, it remains probable when we conditionalize on the antecedent. (Could there be a conditional whose truth-conditions track the proposed assertability-conditions? This is a huge topic about which I have nothing much to say. Suppose such a conditional existed, though, and evidence for p did not carry through to $e \rightarrow p$. This gives us a counterexample to closure only if p implies $e \rightarrow p$, as it presumably doesn't. S&S are aware of this and do not mean to be running the two conditionals together. Some of their formulations, though, may tempt the reader into such a confusion. "Do you know that *If your watch reads 3:00, it is showing the correct time?* Do you know, just by looking at it, that *Even if the watch has stopped, it is showing the correct time?* "Even if" suggests that we retain our confidence in its showing the correct time even on learning that it has stopped.)

1. q is not known *on the basis* of p ; it had to be known beforehand
2. q is not *super*-known; it evokes higher standards than were applied to p
3. one doesn't *know* that q is known
4. the knowledge is *unearned*; we have no real evidence for q
5. epistemic anxiety prevents us from fully *believing* that q

Despite some worthy attempts, however, they have not come up with a unified story about the phenomenon as such. “Hinge” propositions may have to be known beforehand, but there is nothing hinge-y about the dogmatic implication that counterevidence is misleading. I may lose my nerve when it comes to not being a brain in a vat, but I am, I think, pretty confident that I will not be hit by lightning, and that the animal in the cage is not a donkey.

Suppose a better, more unified, account could be found which *vindicated* the plentiful appearances of openness rather than calling them into doubt. That would turn the tables somewhat and put closure fans at a disadvantage. Sharon and Spectre paper seem to have hit on an explanation at the right level of generality. That is what makes their paper so interesting. Closure looks much less defensible if IONs are generated by the mechanism that Sharon and Spectre suggest. How do we decide this issue. The S&S mechanism is explanatory if it is operative when it should be— if it draws the line in the right place. That is what we have been exploring, and will continue to explore for a bit longer.

Recall how it is supposed to be possible to know that Bolt has won the gold, without knowing that he will not be disqualified on the basis of blood tests beyond anyone's current imagination. If Alma emphatically assumes that the blood sample is clean, as (let's say) she is entitled to do, this puts her in a position to learn from blood-neutral evidence that Bolt has won the gold. Some might try to insist that evidence with nothing to offer q could not relieve her ignorance of p . This demand, S&S say, reasonable as it may seem, cannot be pressed in full generality, for it cannot be met in full generality.

The mistake is to believe that one's new evidence together with one's total prior evidence can simultaneously support p and all of its logical implications. But on a probabilistic understanding of the evidence for relation, this is a mathematical impossibility unless p is supported conclusively (that is, with probability 1).

It is mathematically impossible because no matter which e you pick, we can construct from it a consequence q_e of p that e does not support, viz. $e \rightarrow p$.

But, granted that it makes no sense to demand of e that it support every consequence q of p , there is the question of what to conclude from that. To conclude that the demand is generally invalid—inapplicable to any q —seems excessive. The most that follows is that it was overstated. $e \rightarrow p$ is after all a somewhat *outré* example. Couldn't p have a *kind* of consequence that does have to be known for e to close the deal on p ?

Actually, we may have stumbled on one above. It makes no sense that Albert, starting with a clean slate, should learn that Bolt has won gold and

Blake silver from evidence silent on the issue of whether Blake has won silver. This was connected to the fact that *Blake won silver* is part of what is *said* by *Bolt won gold and Blake silver*. (*The wall is red* does not say, even in part, that the wall is red or fails to look red; it doesn't address looks at all.) No counterexamples have yet been given to *immanent closure*:

(IC) If Alma knows that p , and p includes q , that is, has q as a part, then she knows, or is in a position to know, that q .²²

I think we can all agree that parts are subject to more of a closure requirement than "mere consequences." Is there anything in evidentialism to suggest this? Is there anything to indicate why Alma, to know that the wall looks and is red, should know at least that the wall is red? It seems to me that the idea behind S&S's argument works just as well when q is part of p . Imagine for instance that p is $q \& r$. Evidence raising $q \& r$'s probability is as capable of lowering q 's probability as evidence raising p 's probability is of lowering that of $r \supset p$. Evidence raising $q \& r$'s probability can indeed lower the probability of both conjuncts at once. There ought, then, to be cases where Alma knows, say, that it is cold and wet outside without knowing that it is cold outside, or that it is wet. I don't myself know of any.²³ An example closer to home might be this. High rational confidence that I am not a BIV (q) supposedly puts me in a position to learn p , which entails q , from evidence that does not address the issue of vats. This is somewhat believable if p is *I am sitting by the fire*. But p might also be a conjunction with *I am not a BIV* as its first conjunct: *I am neither a BIV nor on fire*. Someone who might, for all they know, be a BIV can hardly claim to know that they are not a BIV and also not on fire. The breakout scenario appears to treat these cases alike.

9 Immanent closure

Immanent closure is meant to be weaker than full closure; that is the whole point. Kripke and Hawthorne show that, if equivalent hypotheses are equi-knowable,

Equivalence (Eq)

If A knows that p , which is (a priori) equivalent to q , A knows that q .

then full closure is recoverable from either of the following:²⁴

Addition (Ad)

If A knows that p , and competently infers $p \vee q$, A knows that $p \vee q$.

Distribution (Di)

²² One should ask in this connection: is it mathematically impossible to probabilify p and all its parts? I would argue that it is not impossible.

²³ Rachael Briggs has some examples that come close ([?]).

²⁴ [?], [?]

If A knows that $p\mathcal{E}q$, then A knows that p and that q .

The proof from (Di): If A knows that p , then by (Eq) she knows that $p\mathcal{E}q$; the two are equivalent given that p implies q . It follows by (Di) that Alma knows that q . The proof from (Ad): If A knows that p , then she is in a position to know that $p\vee q$ by Addition. To know that $p\vee q$ is, however, is to know that q by (Eq); the two are, again, equivalent since p implies q .

Immanent closure had better not endorse these principles lest it blow up into full closure. It endorses Addition only if p says in part that $p\vee q$. But, q may be expected to introduce matters that p says nothing about. *The wall is red if it looks red* (= *The wall is red or does not look red*) addresses itself, for instance, to how the wall looks; it is not included in *The wall is red* because the latter takes no interest in how the wall looks. The principle appealed to here, to which we'll be returning, is

(PA) q is part of p iff the inference p , therefore q is

1. truth-preserving—whenever p is true, q is true, and
2. subject-matter preserving—whatever q is about, p is about.

Distribution is trickier. Immanent closure supports it if $p\mathcal{E}q$ says in part that q —only if *It's a zebra and not a disguised mule*, for instance, says in part that it's not a disguised mule. Perhaps it does in part say this. The inference is certainly truth-preserving. And *It is not a disguised mule* raises, it seems, no issue not raised already by *It's a zebra and not a disguised mule*.

Are we forced then to accept full closure? No, for the objection also relies on (Eq), and (Eq) now looks questionable; it (Eq) obliterates the subject-matter relations that immanent closure feeds on. *It is a zebra* is equivalent in one sense to *It's a zebra and not a disguised mule*; the inference either way is truth-preserving. But the conjunction introduces the matter of mules. That makes it not a part of *It's a zebra* and so not automatically knowable on that basis. Immanent closure does support (Eq) if “equivalence” is taken to involve two-way inclusion. (Eq) on that understanding is harmless, though, since *It's a zebra* does not include the part of $p\&q$ about disguised mules.²⁵

10 Theory sketch

Part-whole is defined in terms of sentential subject matter. Let's first ask what subject matters might be considered as entities in their own right. Lewis suggests they ought at a minimum to determine a relation on worlds: the relation of being alike where the subject matter is concerned.²⁶ He then boldly declares that it *is* the associated relation. A number of properties and relations now become definable.

The state of things in w where m is concerned = the set of worlds m -equivalent to w = the proposition that is true in exactly those worlds.

²⁵ See also [?], [?], and [?].

²⁶ [?]

m is orthogonal to n just if every way matters can stand where m is concerned is compatible with every way they can stand where n is concerned

m includes n just if every way matters can stand n -wise is implied by ways for matters to stand m -wise.²⁷

The number of stars, to use Lewis's example, is the relation of having-as-many-stars-as. How matters stand with it in our world is the set of worlds with exactly as many stars as ours, which for Lewis is the proposition that n stars exist (n being the actual number of stars). The number of stars is orthogonal to the amount of copper insofar as how many stars there are puts no constraints on how much copper there is, or vice versa. It is included in the number of stars and their sizes insofar as answers to *how many stars are there?* are implied by answers to *how many stars are there and how big are they?*

That tells us what subject matters are, and what it is for one subject matter to include another. But nothing has been done to associate a sentence p with *its* subject matter, the p that p is about. This is crucial if we are to judge whether q 's subject matter is contained in p 's, as it must for p to include q .

A truism: whether p is true in a world w supervenes on how matters stand there where its subject matter is concerned. Another thing that supervenes on the state of things where its subject matter is concerned is p 's way of being true: worlds alike where p is concerned cannot differ in how and why p is true. I propose in the spirit of Lewis that p 's subject matter p just *is* the relation that worlds do or don't stand in according to whether p is similarly or differently true in them. Or rather: p 's subject matter is the relation of being true in the same way, and its subject anti-matter \bar{p} is that of being false in the same way. Plugging this into our definitions,

The state of things where where p is concerned in $w = w$'s p -cell (or \bar{p} -cell)
= p 's way of being true in w (or false, as the case may be)

p 's subject (anti-)matter is orthogonal to that of q just if p 's ways of being true (false) are compatible with q 's ways of being true (false)

p 's subject (anti-)matter includes that of q just if every way for q to be true (false) is implied by a way for p to be true (false).

This lets us clarify the definition of content-part:

(PA*) q is part of p just if

1. every world where p is true is a world in which q is true
2. (a) every way for q to be true is implied by a way for p to be true
(b) every way for q to be false is implied by a way for p to be false.

Immanent closure becomes

²⁷ This corresponds more or less to the refinement relation on partitions.

(IC*) If Alma knows that p , then she knows, or is in a position to know, those of its implications that do not “change the subject” —the ones whose subject matter is included in that of p .

Spelled out in terms , Alma is assured of being in a position to know that q if q 's ways of being true (false) are one and all implied by ways for p to be true (false).

11 Saying more

If one looks at a typical list of IONs, all of them have q bringing in subject matter not already there in p . *It's a zebra* is not about painted mules. *I am sitting by the fire* is not about brains or vats. *Bolt won the gold* is not about blood samples. *I turned off the oven* is not about evidence on the question of whether I turned off the oven. *The time is 3am* is not about my watch. The inferences are all *ampliative*, not truth-conditionally, but with respect to their aboutness properties; q introduces additional issues, and in that sense “says more,” than p .

This is not a coincidence. It is its claims about these additional issues, ones that p takes no notice of, that make q harder to know. You might at first think that q would be easier to know than p , since our world is placed in a larger region in logical space. But the shape of the region matters, too. The q -region in cases of interest has, one might say, jagged edges, exposed flanks. The new territory is not as defensible as the old one. Closure violation is the result.

This is all very picturesque, but it's hardly an explanation. An explanation would spell out the mechanism whereby q is newly vulnerable if it is not part of p . We begin with the notion of “a way for q to be true.” Let a *new* way for q to be true is one is not implied by any way for p to be true (likewise for false). What happens when q is not included in p , though p implies it, is that

(NT) q has developed new ways of being true, or

(NF) q has developed new ways of being false.

The question is, why would this make q more epistemically vulnerable?

I am not sure what to say about (NT), but let me take a shot at it. A new way for q to be true is like a new disjunct. A new disjunct is a new opportunity to believe q for the wrong reasons. Alma knows that she turned off the stove by virtue of remembering the fact making it true. It's different with the dogmatic implication that evidence to the contrary is misleading. There were ten witnesses, let's say, and the counterevidence is drawn from their reports. *One* way for q to be true is for the first witness to testify (misleadingly) against you. Another way is for the first two witnesses to testify against you. And so on. Alma presumably thinks that the number is small, since as it grows so does the likelihood she is misremembering.

As we know, though, mistakes on this score can be knowledge-destroying ([?]). One who is right to believe that q is true, but is sufficiently mistaken

about *how* it is true—about how things stand with respect to its subject matter—may not know that q . To know that p you have to be aware of what is going on p -wise, or at least not too confused about what is going on p -wise. Imagine now that Alma, like many people, thinks eye-witness reports largely reliable. Her is keyed, in that case, to one of the earlier disjuncts—she thinks that only one witness testifies against her—when in fact there were three such false witnesses. There is in other words more counterevidence than she imagines, indeed an amount she would find worrisome if she knew of it. Then she for the usual Gettierish reasons may not know that the counterevidence (whatever it may be) is misleading.²⁸ Other examples along these lines could be given.

Now (NF). How a new way of being false could cause trouble is not so mysterious, for these are essentially counterpossibilities, and any number of theories link knowledge to being on top of counterpossibilities. Let the possibilities in question be $\bar{q}_1, \dots, \bar{q}_n$. A sensitivity-type theory might require of Alma that she would have noticed, or had different evidence, had \bar{q}_k obtained. A safety-type theory might see the counterpossibilities as dangerous each in its own way. It could be a probabilistic theory that wants Alma's chance of believing q to be low conditional on any \bar{q}_k . Or an explanatory theory that rules against Alma if the hypothesis that q_k explains how Alma could wind up believing q despite its falsity. Each additional q_k is on any of these accounts one more thing for the would-be knower to be, let us say, on top of. The threat is averted if q is part of p , because then each q_k is a p_k , and Alma had to be on top of the p_k s to count as knowing that p . Parts pick only battles that have already been won.

12 Subject-Matter Sensitivity

Immanent closure is more defensible, it would seem, than full closure. But we need to be careful here. The paradigm of knowing a part is knowing a conjunct; if conjuncts are known, then Distribution holds, which is (almost) all Kripke and Hawthorne need to get full closure back.²⁹

Almost all, because there's also (Eq). This looks initially like the most obvious of the three principles. How can I know just one of a pair of a priori equivalent propositions? It's the same set of worlds, after all; the actual world is in one if and only if it's in the other. Well, a priori equivalent hypotheses may differ in what they're about, and hence in whether their subject matters are included in that of a stronger hypothesis p , and hence in whether they themselves are included in p . If p includes q but not q' , then immanent closure

²⁸ To be clear, this is just an example. I am not saying that Alma *can't* if she plays her cards right know that counterevidence is misleading.

²⁹ I do not say that conjuncts are *always* parts, no matter how logically complex. One may feel that $(p \vee q) \& (p \vee \neg q)$ says the same as p . If so, it cannot contain its first conjunct $p \vee q$, since p does not contain $p \vee q$. No such worries arise with Kripke and Hawthorne's application of the Distribution principle.

requires knowledge only of the first. *I am sitting* is a priori equivalent to *I am sitting and am not a BIV*. But the latter has the larger subject matter, which gives Alma more to be on top of.

Knowledge as we are beginning to conceive it is subject-matter sensitive. This is nothing to run away from; lots of things are subject-matter sensitive: permission, for instance ([?]) and confirmation (see the Appendix). But let's stick to knowledge. Consider the following puzzle. Innocents can be brought in many cases to agree that they might, for all they know, be dreaming. They will then often remark that *their* dreams are, as a matter of fact, not this lifelike. Why do I call this a puzzle? *My dreams are not this lifelike*, which they apparently know, is a priori equivalent to *Nothing this lifelike is a dream (of mine)*, which they apparently don't.

A related example is this. Alma is in fake red barn country. Half of the barn-looking structures are barns and half are fakes; all of the barn-looking objects are red, however. Alma's grandmother calls to ask whether any of the barns she's been perusing are red. Does she know the answer to this? I, at least, am inclined to think that she does. She knows that some of the local barns are red. Granted, not all of the seeming barns are real, but she has laid eyes on dozens of real red barns, and that seems like enough. (Imagine if you like that she has yet to run into a fake barn.) Make of that what you will, certainly Alma comes closer to knowing that some of the barns are red than that some of the local red structures are barns. And yet *Some barns are red* is a priori equivalent to *Some of the red structure are barns*.

This may seem a departure from semantic orthodoxy, but we should not be too scandalized by it. Truth-value shifts brought on by changes in focal stress are old hat in linguistics. (*I always eat TOFU* is a great deal less plausible than *I always EAT tofu*.)³⁰ Changes in the question under discussion can affect truth-value as well ([?]). The focus- and question-sensitivity of knowledge claims in particular has been urged by Dretske and Schaffer ([?], [?], [?]). Knowing that Clyde stole the BEER is not the same as knowing that CLYDE stole the beer. These ideas are in the same ballpark as subject-matter sensitivity.

13 Knowledge by Deduction

Immanent closure applies only when q is contained in p . This may seem at odds with actual epistemological practice. Drawing consequences is a way of extending knowledge; it is not as though we stop to check that the consequence is a part. One tempting idea here is that the consequence *becomes* a part when it is reached, due to retroactive adjustments in what is said by the premise. To see what is going on here, we have to consider the subject-matter sensitivity of knowledge in light of the context-sensitivity of subject matter, in particular, its sensitivity to discourse context.

³⁰ [?]

So, for instance, *The coach is in money trouble* is prima facie about how matters stand with the coach. If the issue is who to approach with a bribe, however, the focus shifts to money trouble and who is it. Here the relevant context takes place before the statement whose subject matter is in question, but there is nothing to stop us from looking also at what comes after. This includes, I suggest, the uses to which p is put, in particular the *deductive* uses.

All of us know the cognitive switch that is pulled when *Numbers exist* is inferred from some humdrum arithmetical claim that no one would think to doubt. *There are primes over ten* seems in some way to *say* more than it did on the math test. *I have a hand* takes on a bold new aspect when Moore uses it to refute the skeptic.

One sort of contextualist will say that the truth-conditional content grows, as the context set expands to include weirder worlds. I am not sure this is right; anyway I want to try something different. Rather than p 's truth-conditions ramping up when q is deduced from it, its subject matter expands, as the conclusion's ways of being false are taken on by the premise. That p acquires new ways of being false (in the sense that the $\neg p$ -worlds are further divided up) may necessitate a reconsideration of that premise.³¹

14 Skeptical Deductions

Let $p+$ be p with the puffed up subject matter; the BIV worlds, for instance, are singled out for attention by the subject anti-matter of *I am sitting by the fire*. $p+$ is true/false in the same worlds as p , but for different (no less specific) reasons. $p+$ contains q where p did not. Two cases can be distinguished: we know p but not $p+$; we know $p+$ in addition to p .

Suppose that we know $p+$ in addition. Then since q is contained in $p+$, we know that q , by immanent closure. This is the ordinary case, in which knowledge is preserved under deduction. Knowing that I am sitting gives me knowledge that I am standing or sitting, because I continue to know I am sitting, even when its subject anti-matter expands to include neither-standing-nor-sitting as a distinguished alternative. To put it in Nozickian terms, I would have noticed, had I been neither standing nor sitting.

Imagine now that p is known but $p+$ is not. That knowledge is lost when p is expanded to include q suggests that q is contributing counterpossibilities that are not as easy to pick up on. I might infer that I am not a BIV from the known fact that I am sitting. *I am sitting* then undergoes topical expansion, with being-a-BIV now a separately articulated counterpossibility. To know what *I am sitting* NOW says, I must know that I am not a BIV. Since I do *not* know this, so the story goes, I cannot be said any longer to know that I am sitting. Using "thought" for propositions with subject matter included, I know the thought that *I am sitting* used to express, but not the one it expresses now.

³¹ Imagine a tourist map of Bel Air, showing where the stars live, is produced in a legal dispute about oil rights.

In regular contexts, I know that I am sitting. In skeptical contexts, I supposedly do not. This cannot not be the whole story, for lightweight propositions seem better known than their heavyweight consequences even in philosophy class. Even there, I feel in a better position with respect to *I am sitting* than with respect to *I am not a BIV*. Standard-issue contextualists have trouble explaining this. Neither proposition is known, judging by higher standards, both are, judging by lower.

The difference is that lightweight propositions retain *even in skeptical contexts* a substantial known part. Let the original, pre-skeptical subject matter of *I am sitting* be Yablo's posture. The alternatives recognized by Yablo's posture, are, say, I am standing, I am lying down, I am leaning, I am hanging by my heels, and (a final catch-all category) "other." The part of *I am sitting* that concerns my posture is still kicking around inside, even when the sentence has come to say more.³²

Seeming closure violations have been met with three main responses: counterfactualism (Nozick), contextualism (Cohen, DeRose, Lewis), and Carnap's idea that *There are numbers* seems harder to know than *The # of martian moons = 2* because we hear it as making an "external" claim. Our picture has some contextualism in it, since the subject matter of *I am sitting* changes in skeptical contexts, thereby "destroying our knowledge." It has *potentially* some counterfactualism in it, if being on top a counterpossibility is being such that that one would have noticed, had that counterpossibility obtained. It has some Carnap in it, too, for we can always, when the doubters come round, take refuge in the ordinary, "internal," part of *I am standing*, the part that concerns its old, non-skeptical, subject matter. The ordinary part we do know.

Appendix: Hempel was Right³³

The golden age of confirmation theory³⁴ begins with Hempel's "Studies in the Logic of Confirmation." He articulates there four possible conditions on evidential support.

ENTAILMENT: *e* confirms any *h* that it entails.

CONSISTENCY: If *e* confirms *h*, it does not confirm anything contradicting *h*.

SPECIAL CONSEQUENCE: If *e* confirms *h*, it confirms what *h* implies.

CONVERSE CONSEQUENCE: If *e* confirms *h*, it confirms that which implies *h*.

A fifth principle, mentioned in passing, is

CONVERSE ENTAILMENT: *e* confirms any *h* that entails *e*.

Hempel accepts the first three of his proposed conditions, but not the two CONVERSES. What mainly bothers him about them is that they (either one)

³² See [?] for "the part of *p* about *m*." It's a proposition that is false in *w* just if *p* is false there for reasons visible to *m*. The part of *I am sitting* that concerns my posture is false if I am standing or lying down, but not if I am a BIV.

³³ Much of what goes on in this appendix is due to Ken Gemes ([?] [?] [?] [?])

³⁴ [?][?] [?] [?], [?]

trivialize the confirmation relation, given ENTAILMENT and SPECIAL CONSEQUENCE. For let e and f be arbitrarily chosen.

1. e confirms e (ENTAILMENT)
2. e confirms $e \& f$ ((1), CONVERSE CONSEQUENCE)
3. e confirms f ((2), SPECIAL CONSEQUENCE)

This objection has been found puzzling. Why put the blame on CONVERSE CONSEQUENCE? Its contribution is only to get us to (2): e confirms $e \& f$. But (2) is an instance of CONVERSE ENTAILMENT, which is apt to seem obvious.. If h entails e , so that $\neg e$ precludes h , how can e not speak in favor of h , if only by removing a possible obstacle? CONVERSE ENTAILMENT is backed, too, by the Bayesian analysis of confirmation: $\text{pr}(h|e)$ exceeds $\text{pr}(h)$ if h entails e . SPECIAL CONSEQUENCE, on the other hand, is from a Bayesian perspective untenable. Evidence making h likelier *cannot* make all its consequences likelier, and there are particular consequences whose probability is bound to go down.

How could Hempel have been so wrong? Why will he not let go of SPECIAL CONSEQUENCE, when the problems seem obvious? Carnap suggested a diagnosis. He thought Hempel was mixing up two notions of confirmation, whose differences emerge when we consider the matter quantitatively. Let $c(h, e)$ be the degree to which e confirms h . e confirms h *incrementally* if $c(h, e \& k) > c(h, k)$, for k some body of background information. It confirms h *absolutely* if $c(h, e \& k)$ exceeds some chosen parameter, let's say .98.

Absolute and incremental confirmation should definitely not be confused. But is Hempel confusing them? One would expect Carnap to argue that some of Hempel's conditions hold for the absolute notion, others for the incremental notion. But all of Hempel's preferred conditions hold for the absolute notion! It is only CONVERSE CONSEQUENCE, which he rejects, that fails to hold absolutely.

The problem is that Hempel's *rhetoric* and his *examples*, which mostly involve the confirmation of generalizations by their instances, suggest the relative notion. A black raven makes it likelier, not absolutely likely, that all ravens are black. Relative confirmation is naturally understood as positive probabilistic relevance, or probabilification. And probabilification meets neither of Hempel's two main conditions.³⁵ Not CONSISTENCY, for *Rudy is a black raven* is positively relevant both to *Rudy is a happy raven* and *Rudy is an unhappy raven*. Not SPECIAL CONSEQUENCE, for *Rudy is a black raven* relatively confirms *Rudy is a black raven and all other ravens are white* despite being negatively relevant to *All other ravens are white*.

At this point Hempel might seem refuted. His conditions hold for absolute confirmation, but that is not what he's talking about. The standard (Bayesian) model of relative confirmation is positive probabilistic relevance, but that interpretation does not meet his conditions. This doesn't entirely settle the matter, however, for a reason noted by Earman:

³⁵ It meets only ENTAILMENT.

...there may be some third probabilistic [notion of] confirmation that allows Hempel...to pass between the horns of this dilemma...it is up to the defender of Hempel's instance-confirmation to produce the *tertium quid* ([?]. p. 67).

Hempel in fact left a number of clues about this. Here he is introducing the stronger condition of which SPECIAL CONSEQUENCE is meant to be a corollary:

an observation report which confirms certain hypotheses would invariably be qualified as confirming any consequence of those hypotheses. Indeed: *any such consequence is but an assertion of all or part of the combined content of the original hypotheses* and has therefore to be regarded as confirmed by any evidence which confirms the original hypotheses. This suggests the following condition of adequacy:

GENERAL CONSEQUENCE CONDITION (GC): If an observation report e confirms every one of a class P of sentences, then it also confirms any sentence $[q]$ which is a logical consequence of P ([?], 103, italics mine)

Hempel's reasoning here is interesting. P 's consequences are confirmed by e , he says, because "any such consequence is but an assertion of all or part of the combined content of the original hypotheses." Supposing for simplicity's sake that $P = \{p_1, p_2\}$, Hempel thinks that any consequence of $p_1 \& p_2$ asserts part or all of the combined content of p_1 and p_2 , and that this helps us to see why e 's support for p_1 and p_2 would carry through to q .

Why does Hempel insist on e confirming "every one" of the sentences in $\{p_1, p_2\}$, as opposed, say, to either of them, or their conjunction? If one says *either*, then e confirms any f that you like, by virtue of confirming a member (the first) of $\{e, f\}$. Similar difficulties arise if it is the conjunction we focus on; f might be a free rider in $e \& f$. The point is that GENERAL CONSEQUENCE would not be plausible, if e were not asked to confirm each of P 's members separately. Let's hold onto this as it will be important later.

Given that Hempel insists in GENERAL CONSEQUENCE on "wholly" confirming evidence—evidence confirming *both* of p_1, p_2 —why does he not also insist on wholly confirming evidence in SPECIAL CONSEQUENCE—evidence confirming each conjunct? Any reason there might be for asking e to confirm both members of $\{p_1, p_2\}$ is surely also a reason for asking it to confirm both conjuncts of $p_1 \& p_2$! SPECIAL CONSEQUENCE as we read it today imposes no such requirement, which to me suggests that we may be reading it incorrectly.

Again, Hempel objects to CONVERSE CONSEQUENCE that *Rudy is black* does not confirm *Rudy is black & Hooke's law holds*. But, Rudy's blackness *does* confirm the conjunction in the sense of probabilifying it. What it doesn't do is confirm all of the conjunction; Hooke's Law is not made the least bit likelier. Charity requires us to interpret him as imposing the stronger requirement: To confirm a conjunction, e must confirm both conjuncts. Bayesian confirmation lacks this property, but we can easily impose it:

(WC) e wholly or pervasively confirms h iff e probabilifies all of h 's parts.³⁶

³⁶ That is, $\text{pr}(g | e \& k) > \text{pr}(g | k)$ for each $g \leq h$.

Wholly confirming h is confirming *all* of it. Let's review Hempel's conditions with this notion of confirmation in mind.

CONSISTENCY: *Rudy is black and happy* is incompatible with *Rudy is black and unhappy*, and yet both are probabilified by *Rudy is black*. The problem here is non-pervasiveness. To probabilify both statements *and their parts*, *Rudy is black* would have to make it likelier both that Rudy is happy and that Rudy is not happy.

ENTAILMENT: Suppose e entails h . Then if g is part of h , e entails g by transitivity of entailment, whence $\text{pr}(g|e \& k) = 1$. e pervasively confirms h , then, provided only that $\text{pr}(g|k) < 1$. Call a hypothesis h *novel* if it has no parts g such that $\text{pr}(g|k) = 1$. Evidence e that entails a *novel* hypothesis confirms *all* of that hypothesis.

SPECIAL CONSEQUENCE: If e confirms all of h , then it confirms all of h 's parts, and hence (by transitivity of part-whole) the parts of its parts. To probabilify all of the parts' parts is the same as pervasively probabilifying each part. Let HEMPELIZED SPECIAL CONSEQUENCE be the principle

(SC _{h}) If e pervasively confirms h , then it pervasively confirms h 's parts.

This is virtually a logical truth! Hempel's version of SPECIAL CONSEQUENCE has its problems, but implausibility is not one of them; triviality is more like it.

A word finally about Hempel's positive theory. Though rejecting CONVERSE CONSEQUENCE, he thinks a certain *kind* of consequence *is* confirming. A generalization h is confirmed by its "development" for observed individuals, written $Dev_I(h)$. Now, h 's development for I sounds like it should be the part of h that concerns I , that is, what h says about those particular individuals. Hempel's positive theory would then be that a generalization is confirmed by its parts.

If that is the intention, though, the definition doesn't capture it. $Dev_I(h)$ is defined by Hempel as what h says when its quantifiers are restricted to the individuals in I . This is not always even a consequence of h , let alone a part.

To see the problem, let pluralism be the theory that for all x , there exists a y that is not identical to x . Pluralism is true, I take it, hence true about every subject matter. But its development for I is false, if I contains just one individual. There is a problem in the other direction too. Let monism be the theory that for all x and y , x is identical to y . Monism is false, I take it, and there is not a lot of inductive evidence for it. Hempel has us gaining such evidence whenever we refuse to look at more than one individual. (There is indeed just one thing, leaving aside all the other things.)

One might think of Hempel as groping here for the notion of partial truth—truth where a certain subject matter is concerned. For h to be true about subject matter m , recall, corresponds to the truth outright of the part of h that concerns m . The problem we were running into above is that the part of h about m cannot always be obtained by restricting quantifiers. This is a very small mistake on Hempel's part! Yet it's the only one he has made, on the present interpretation.