

The truth and something but the truth

Stephen Yablo

The philosopher Carl Hempel, in whose honor these lectures are given, tells somewhere of some other lectures, on the topic of material objects, given by Rudolf Carnap at Harvard in the 1930s. Carnap is supposed to have begun these Harvard lectures by saying,

"Let A be some physical body, such as a stone, or a tree, or – to borrow an example from Russell – a dog."

I wish I could explain my topic the way Carnap explained his: with examples. But I am going to be talking about aboutness, subject matters, meaning, truth, reasons for truth, contents, parts of contents – as to be female is part of being a vixen – extricability of – femalehood is more extricable from vixenhood than red from scarlet – and philosophical applications of the above. These sorts of notions don't especially lend themselves to introduction by example, or to the extent they do, the examples won't mean much except surrounded by so much philosophical commentary as to defeat the purpose. Perhaps though I can set the mood with some stories.

The first story, which is true, concerns my wife Sally. Sally's dissertation was on the same sort of topic as Carnap's lectures: physical objects and their identity over time [insert Sally story here--gist is she tells her interviewers she will be speaking of properties as a convenience but the issue is really material objects and how they change. "It remains the same tomato when it *acquires the property of redness*" vs "It remains the same tomato *when it turns red.*"....]

The aspect of the story that interests me is this. Sally made a statement implying the existence of properties, one that she knew was false if properties don't exist. But she was absolutely unbothered by the possibility that properties don't exist. Her excuse for this insouciant attitude was that her topic was material objects, not the properties if such there be of those objects. Which raises an interesting question: *How on earth it is an excuse for speaking falsely that one's topic was so and so?* How is speaking falsely about a particular topic any better than speaking falsely with topic unspecified? The answer to that is suggested by my second story; yes Sally may have been speaking falsely, and yes her topic was the tomato, but that doesn't mean she was speaking falsely about the tomato.

The second story is due to Nelson Goodman and Joseph Ullian, in a paper called "Truth about Jones." Jones is on trial for murder and Falstaff is the chief witness for the defense. Jones's attorney concedes there is a problem with Falstaff's testimony. It is false. That would seem to be a fairly serious problem, but the

attorney (Lupoli, he's called) thinks he sees a way out. The testimony was about his client Jones, no getting around that. And it was false, no getting around that either. But, Lupoli insists, the testimony was not *false about Jones*. The judge rejects this as nonsense and shuts down the proceedings, threatening Lupoli with contempt unless he can explain what he could possibly have meant by false and about Jones yet not false about Jones.

I hope you see a connection with the Sally story. Just as Sally wasn't concerned if her statement was false, Lupoli doesn't care if Falstaff's testimony was false. It's enough for Lupoli if the testimony was partly true – true in what it said about Jones. Maybe that should be Sally's excuse too; it's enough for her if her statement was partly true, true in what it said about the tomato.

Now this idea of being partly true, and the related idea that partial truth might be quite enough in certain cases, does not come naturally to us. It is hard not to share the judge's frustration when he threatens Lupoli with contempt. The phrase "partly true" is good English, of course. Apparently it was decent Greek too, for the creation myth in the Phaedrus is described by Socrates as "partly true and tolerably credible." And when Cratylus complains it would be nonsensical to address Socrates using somebody else's name, Socrates says, "Well, but [it] will be quite enough for me, if you will tell me whether the nonsense would be true or false, or partly true and partly false: - which is all that I want to know." That is actually not a bad statement of one theme of these lectures: sometimes whether a statement is partly true is all that we want to know.

Why then do I say that it doesn't come naturally to us to settle for partial truth? Consider a third story, this one by the psychoanalyst Melanie Klein. (She didn't consider it just a story of course.) Newborns in Klein's view face an enormous challenge, that of putting the things that gratify them together with the things that frustrate them into a single world; they come, as Klein would put it, to appreciate that the good breast and the bad breast *are the very same breast*. This cognitive hurdle is usually cleared at around 4 months; that's when the infant moves from the paranoid-schizoid position to the apparently far preferable depressive position.

That's the normal case. Occasionally though the integration challenge proves too great, and the individual never really wraps their mind around the fact that one and the same thing can be partly good and partly bad. The result is the cognitive style known as black/white thinking or polarized thinking. A black/white thinker is the type of person who loves you or who hates you, flipping back and forth according to how recently you've disappointed them. More generally they're the type of person who insists on dividing the world up into good, full stop, and bad, full stop. This kind of attitude is familiar with kids,

of course, and more forgivable there. I remember my Isaac squirming around in his seat at the movie Shrek: he couldn't relax until he knew whether Donkey (the Eddie Murphy character) was a good donkey or a bad donkey. No problem there. But imagine you're watching the news with an adult and all they want to talk about is, Is this George W. Bush fellow a good man or a bad man? When you try to suggest it's more complicated than that, he is good in these ways and not so good in those others, they reject this as spineless evasion. Answer the question, they say. That's black/white thinking and it surely deserves its reputation as pathological.

But our assessment changes when the focus shifts from goodness to truth. Demanding to know whether a statement is *true*, full stop, or false, full stop is not pathological at all; on the contrary it's considered forthright and healthy minded. It's almost as if having lost our Kleinian paranoia about goodness, there was no energy left to outgrow the analogous attitude about truth. A second theme of these lectures is that this is nevertheless worth doing, or insofar as we've already done it, owning up to doing. Let's move beyond the paranoid/schizoid position about truth and move boldly ahead into the, well, I guess it's the depressive position. (I admit it's not the best rallying cry.)

There are two questions at this point: what is partial truth? and why should we put up with it? The second question I want to leave to the end, but the quick answer (adapted from the blues classic "Born Under a Bad Sign": "if it wasn't for bad luck, I wouldn't have no luck at all") is that there are certain areas where if it wasn't for partial truth we wouldn't, or might not, have any truth at all. But as I say, I propose to leave to the end. For now I want to focus on the other question: what is partial truth in the first place?

Here is the naivest idea possible about partial truth. I use the word "hypothesis" to stand ambiguously for a sentence or its content:

(1) A hypothesis is partly true to the extent that it has parts that are (wholly) true.

But what is meant by part of a hypothesis? The naivest possible idea about that is

(2) One hypothesis is part of, or included in, another iff it is implied by the other.

(*A includes B* iff *A implies B*.)¹ The naivest possible idea about partial truth is correct, I think; something is partly true to the extent it has parts that are wholly true. But the naivest possible idea about parthood (inclusion) is questionable.

A paradigm of inclusion, I take it, is the relation that typical conjunctions bear to their conjuncts – the relation *snow is white and grass is green* bears, for example, to *snow is white*. A paradigm of non-inclusion is the relation disjuncts bear to disjunctions; *snow is white* does not have *snow is white or grass is green* as a part. This is despite the fact that disjuncts imply disjunctions every bit as much as conjunctions imply conjuncts. On the face of it, then, there is more to inclusion than implication.

You might say that paradigm case intuitions are a poor basis for theory, and I might agree with you. But we're dealing with more than one-off intuitions here. The intuitions are quite systematic, and they suggest that parthood has an explanatory role to play that requires it to be more than mere implication.

Saying: Someone who says that snow is white and grass is green has among other things said that snow is white. But to say that snow is white is not *inter alia* to say that it is white or grass is green. Why, when there is implication in both cases? Saying-that transmits down to the parts of what is said but not to mere consequences, meaning by this consequences that aren't parts.

Permission: Permitting someone to eat pork chops is, other things equal, permitting them to eat pork. Permitting them to eat pork, though, is not permitting them to eat pork or human flesh. Why this difference, when eating pork or human flesh is every bit as much implied by eating pork as eating pork is by eating pork chops? Maybe because one permits (normally) the parts of what one permits, but not implications more generally.

Priority: Conjuncts seem somehow prior to – preconditions of – their conjunctions, while disjunctions are posterior to – consequent on – their disjuncts. p must hold before $p \& q$ can hold, but it is not the case that

¹ Consider the version of (2) that focuses on contents rather than sentences. If we think of contents the way Lewis does, as sets of possible worlds, then (2) says one hypothesis includes another iff it is a subset of the other. Lewis in effect proposes this at one point, in "Statements Partly About Observation." He also notices that it doesn't sit very well with his view that subsets are parts of the sets they sub – which leads him to define contents for these purposes as the set of worlds ruled out rather than in. I don't want to dwell on these issues since (2) strikes me as pretty clearly mistaken.

$p \vee q$ must hold before p can hold. It is tempting to reach here for a more general principle: parts are prior to wholes, and other consequences – mere consequences – are posterior to, or logically downwind from, their impliers.

Explanation: The falsity of a conjunct p explains the falsity of the conjunction $p \& q$. But the falsity of a disjunction $p \vee q$ only guarantees, without explaining, the falsity of a disjunct p . Why? It makes sense that if S has a false part, S will be false thanks to the falsity of that part. A part's falsity is well-positioned to explain the falsity of the whole. But that S implies a random falsehood seems more a symptom of S 's falsity than the reason for it.

And then of course

Partial truth: *Snow is white and grass is blue* is made partly true by the fact that *snow is white*. *Grass is blue* is not made partly true by the fact that *snow is white grass is blue*. Why not? True parts confer partial truth on wholes. Other true implications have no such power.

So more is involved in B 's being part of A than B being implied by A ; parts are special and behave differently than mere consequences. What is the missing ingredient? What is the X in

Parthood = implication + X ?

The answer to this is implicit in the stories we began with. Falstaff's testimony is partly true because *the part that concerned Jones* is true. Sally's interview statement is partly true because *the part that concerns the tomato* is true. A statement's parts are identified in both cases by looking for an implication *whose subject matter is part of the subject matter of the original statement*. This suggests that for B 's content to be part of A 's involves, in addition to A implying B , that *B 's subject matter is part of A 's subject matter*. Conversely the reason A 's implications are sometimes *not* parts of A is that they bring in new subject matter. *Snow is cold* does not include *snow is cold or grass is green* because the latter brings in the matter of grass's color which is wholly absent from *snow is cold*.

So: contents have parts, but to identify them we will have to broaden our focus from truth-conditions to what sentences are about, their subject matters. A third theme of these lectures is that there are lots of things to understand which it helps to broaden our focus from truth-conditions to subject matters.

“Broaden our focus” suggests subject matter has been relatively neglected in philosophy. This is true, I think. How many times have you heard a philosopher question the analysis of X as Y on the grounds that their truth-conditions differ; X can be true when Y is false or vice versa? (Let’s see, there was Plato on justice, Gettier on knowledge, Frankfurt on free action, ...) If one is serious about subject matters, there’s a completely different way to challenge the proposed analysis of X as Y: Y gets the subject matter wrong; X is about one thing Y is about something else. How many times have you heard a philosopher argue like that?

Subject matters have been relatively neglected, not completely neglected. An example of non-neglect is Frege in his work on the informativeness of identity statements. Frege initially held that "Hesperus = Phosphorus" says of the words "Hesperus" and "Phosphorus" that they refer to the same object. His subsequent objection to this account was not that it assigns the wrong truth-values, but that it gets the subject matter wrong: "H = P" is about the planets Hesperus and Phosphorus, not our devices for picking those planets out.

Or consider Kripke’s famous objection to counterpart theory. Humphrey is upset that he could have won the election but didn’t. Counterpart theory understands his possibly winning the election as the winner being in some other world a guy that suitably resembles Humphrey. But "Humphrey could care less whether someone else...would have been victorious in [another] world" (45). This has been called the argument from concern, but the lack of concern stems from a prior problem: "if we say, 'Humphrey might have won the election'....we are not [according to counterpart theory] talking about something that might have happened to *Humphrey* but to someone else, a "counterpart"" (45). Humphrey doesn’t care about what the counterpart-theorist is offering him because it has the wrong subject matter; it’s in part about somebody else.

One might even argue that Kripke gets the subject matter wrong in places. Consider how he explains the intuition that heat might have been low molecular energy; we confuse *that* possibility with the possibility that low molecular energy could have felt this way to people with different neural wiring than ours. But, one might say, the intuition that *this* could have been low energy is an intuition about *heat*, not the way heat feels to local observers whoever they may be. The possibility Kripke points to *might* explain the intuition that low molecular energy could have felt this way *to us*; for we could see the other-worldly observers as our epistemic counterparts. But the intuition that *that* could have been low

molecular energy is not an intuition about us; it's about a switcheroo out there, where the heat is, not back here where we are.²

One last example, this one concerning epistemic modality instead of metaphysical. The standard analysis of "It might be that φ " has it expressing something in the neighborhood of "I don't know that not- φ ." But that can't be right. "Bob might be in his office" is about Bob and his office, not the extent of the speaker's knowledge. (A: "Bob might be in his office." B: "Will you please get over yourself?") There is a concern aspect here too. That Obama is might win the election is presumably a source of great satisfaction to him. But the limited extent of his knowledge doesn't excite him one bit.³

But, these sorts of examples notwithstanding, subject matter is usually treated as just a way station on the road to truth-conditions. And sometimes a relatively *arbitrary* way station at that, because, as many authors have been concerned to emphasize, one can scramble what subsentential expressions refer to while leaving truth-conditions the same; this is argued in various papers of Quine's, Davidson's "Reality Without Reference," and Putnam's "Models and Reality." Quine seems to suggest that not much would be lost if we assigned numbers to every material object and read statements seemingly about them as really about the associated numbers; "I am hungry" would say of the number 18 that it's hungry*, where to be hungry* is to have your associated person be hungry.

Truth-conditions underdetermine subject matter; that's true. One could conclude that subject matter is less important or even somehow less real. But the more natural conclusion is that we have in subject matter a potentially independent factor in overall meaning, one that can vary even as truth-conditions remain the same.

A good, anyway tempting in the present context, illustration is Hempel's raven paradox. ⁴ The paradox raises lots of questions but one is this: *All ravens are*

² A proper Kripke-style explanation of the intuition that *that* (heat) could have been low molecular energy would point to a doppelganger of heat with low molecular energy as its underlying constitution. It would point to the possibility of substituting low energy for high while retaining outward appearances. There is no such possibility, however, which is presumably why Kripke posits a switcheroo on the observer side rather than the observed.

³ Seth Yalcin makes a related point.

⁴ A simpler example. *Some barns are red* is truth-conditionally equivalent to *Some red things are barns*. But the one is about the color of barns, the other is about the structure-type of red things. As a result they seem to say different things. Travellers through fake barn country might know the first but not the second.

black and *All non-black things are non-ravens* are truth-conditionally equivalent. One would expect, then, that evidence confirming the one should equally confirm the other. And yet a black raven seems more confirmationally relevant to *All ravens are black* than *All non-black things are non-ravens*, while a non-black non-raven seems (if anything) more confirmationally relevant to *All non-black things are non-ravens* than to *All ravens are black*.

Here is what I think we are initially tempted to say, before the paradox experts talk us out of it. *All Fs are Gs* and its contrapositive may be true in the same worlds. But they are about different things. One is about Fs and whether they're G; the other is about non-Gs and whether they're F. It seems only natural that to confirm a hypothesis about one kind of thing, one should look at examples of that kind of thing, while to confirm hypotheses about another kind of thing, one should look at examples of that other kind of thing.

This is not a very sophisticated reply but I am tempted to think it's correct. All it's missing is some kind of story about how a difference in subject matter (with truth-conditional content held fixed) can make for a confirmational difference. How about the following: Already in Goodman we find the distinction between an instance confirming a generalization in the sense just of eliminating one potential counterexample--the way a coin coming up heads once confirms that it will come up heads ten times in a row--and the kind of confirmation that's supposed to occur in induction, where this F that is G makes it likelier that other Fs not yet observed are G. For an instance to inductively confirm a hypothesis it has to probabilify all the hypotheses' other instances.

But an instance of a generalization is surely a *part* of that generalization. So to inductively confirm a hypothesis (in this case anyway) is to probabilify the hypothesis *and its various parts*. A change in parts can thus make for a change in what inductively confirms a hypothesis. Recall now that parthood is defined in terms of subject matter. This completes the story: subject matter differences bring about mereological differences bring about confirmational differences. Truth-conditionally equivalent hypotheses that differ in what they're about, and hence in what other hypotheses they have as parts, are going to be inductively confirmed by different evidence (see also the appendix).

So that's the fourth theme: subject matters are an important and potentially independent factor in meaning, over and above truth-conditional content. The fifth theme is that they are not as independent a factor as all that, because our notion of subject matter is deeply tied up with our notion of why a sentence is true. To say that *All ravens are black* is about *ravens* and not pencils is very like saying that it is true because of what *ravens* are like and not because of what pencils are like. What are the underlying principles that link subject matter, on the one hand, to the reasons for a sentence's truth-value, on the other?

There are two, I think, and they're both to do with *changes* in subject matter vs changes in why a sentence is true. I should explain that by changes in a sentence's subject matter, I don't mean changes in the *identity* of its subject matter. Take *The President's father was (also) a politician*. I assume that this, has the numerically same subject matter in w_1 , where the President is Bush, as in w_2 , where the President is Gore. But as we know, numerical identity is compatible with qualitative distinctiveness. That numerically unchanged subject matter takes on a different shape, or is differently realized, when the President switches to Gore. Suppose I'm a reporter on the *President-is-son-of-a-politician* beat, and my editor calls to ask if anything's changed. I'm going to say No if the President owns an additional bicycle, and Yes if the Presidency goes to Gore. A sentence's subject matter changes, as I'm thinking of it, if it changes qualitatively in the manner just imperfectly indicated, not numerically as it might if the sentence took on a whole new meaning.

I said there were two principles linking a sentence's subject matter to the reasons for its truth-value. One principle is that if a sentence's subject matter is unchanged between worlds, then the sentence's truth-value cannot change, *nor can the reasons for its truth-value change*. Where reasons for truth or falsity change, subject matter cannot remain the same. Consider again $A =$ *The President is the son of a politician*. If I tell my transworld editor that A is true in w_1 because George W Bush's father is George HW Bush, and in w_2 because Al Gore's father is Albert Arnold Gore, Sr, then I can't go on to say that there's nothing to report, A 's subject matter is entirely unchanged. Reasons for truth cannot change without subject matter changing too. Subject matter is at least as fine-grained as reasons for truth.

I would argue conversely that subject matter is no more fine-grained than reasons for truth. Suppose that Bush owns one more bicycle in w_3 ; otherwise things are pretty much the same, in particular, Bush is President in both and has the same father with the same career in both. Then it would seem that w_1 and w_3 are just alike with respect to the subject matter of A . If we had to say *why* changing the number of bicycles isn't mucking with that sentence's subject matter, the answer would be that *how many bicycles Bush owns is irrelevant to why the sentence is true*. All that matters to why the sentence is true is the President and his father and his father's career; the number of bicycles is neither here nor there. This is only one example but it suggests that worlds that differ in ways irrelevant to why A is true do not differ at all with respect to A 's subject matter. The relation between truth-conditions and subject matter therefore looks to be this: whereas we give A 's *truth-conditions* by saying of each world *whether* A is true in that world, we give its *subject matter* by saying *why* it is true in that world.

This identification of subject matter with reasons for truth might lead one to wonder whether subject matter is indeed a further factor in meaning, something not automatically recoverable from a sentence's truth-conditional content (= the set of worlds where it is true). I will grant you that specifying in which worlds *A* is true may tell us quite a lot about why *A* is true in those worlds; it may even tell us all there is to know. Example: A disjunctive proposition like *x is round or y is square* will have as its truth-conditional content the union of two sets of worlds both of which are much more natural than it. To the extent that reasons for a sentence's truth are natural, proportional, guarantors of its truth, *A*'s truthmakers will presumably be, first, that *x* is round, and, second, that *y* is square

But to say that truth-conditional content sometimes determines subject matter and always constrains it is not to say that it always determines it. Think again of the raven paradox. *All ravens are black* and its contrapositive are true in the same worlds but for different reasons: the one because of what the ravens are like, the other because of what a completely disjoint population is like, viz. the non-black things. Another example to which we might or might not be returning later: consider *The king of France is bald* and *There is a bald person who is identical to the king of France*. These are logically equivalent so presumably true (false) in the same worlds. Why then does only the second *strike* us as false. Well, the first is about French kings and false because there are none; the second is about bald people and false because none of them is a French king. Sentences that are false *just because their presuppositions are false* strike us as not getting out of the starting gate; a sentence has to make a second mistake, beyond the failure of its presupposition, before it will strike us as false. Thus a shift purely in subject matter can change a sentence's perceived truth-value, in this case from non-false to false.

I have been speaking of reasons for truth, sometimes called truthmakers. But I mean a different thing by this phrase than some people. Truthmakers are often seen as entities of a different ontological category than what they make true: they are, say, objects or tropes making true sentences or propositions. That's not the view taken here. Truthmakers are facts, conceived as true propositions; they are thus of the same category as at least some of the things they make true, viz. other true propositions. Another difference is that to make a proposition *P* true involves more than just necessitating its truth. *P*'s truthmakers should tell us *why P* is true, in the same sort of way as an effect's causes tell us why it happened. *T* is *P*'s truthmaker only if *P* is true because of *T* or in virtue of the fact that *T*.

T makes *P* true in *w* iff

- (i) *T* is true in *w*,
- (ii) *T* implies *P*, and

- (iii) P is true because of T or in virtue of the fact that T.

(This is more or less what Armstrong says too.)

What about the in-virtue-of relation? I have no real theory to offer, but it seems to me that truthmakerhood typically involves a tradeoff between two constraints. The first is that truthmakers should be *natural*, both in an extensional and intensional sense. A fact is extensionally natural to the extent that going on in a smaller, more connected, spatiotemporal region guarantee that it obtains. A fact is intensionally natural to the extent that it obtains in a more natural set of possible worlds. The other constraint is *proportionality*, which again has two components. Truthmakers should be *efficient*: T should not be weakenable to another P-implier that is no less natural. (What makes it true that **There are cats**? Efficiency favors **Yossele is a cat** over **Yossele is a vegetarian cat**.) Truthmakers should be *robust*: if P holds in a bunch of different sorts of worlds, T should as far as possible make it true in worlds of each sort. (What makes it true that **A raven is flying** \supset **something black is flying**? Robustness favors **Ravens are black** over **Something black is flying**.) For T to make P true is *something in the neighborhood* of T effecting the best possible tradeoff between the demands of naturalness and proportionality.

A truthmaker for P in world w is (roughly!) a fact that

- (i) obtains in w,
- (ii) implies P, and
- (iii) is as natural and proportional to P as a fact satisfying (i) and (ii) can be.

This brings us to the notion of sentential subject matter, which I have said is related to the reasons a sentence (or content) is true.

David Lewis defined a subject matter as a partition of logical space – a division of all possible worlds into jointly exhaustive, pairwise disjoint, subsets, called “cells.” The subject matter **the number of stars** partitions logical space into the worlds where there are no stars, the worlds where there are one star, and so on for each possible number of stars. Bearing in mind that sets of worlds for Lewis are propositions, a subject matter can also be conceived as a set of propositions, specifying between them the various ways things can be where that subject matter is concerned. So, **the number of stars** is also a set of propositions, each to the effect that **there are exactly k stars** for some specific value of k.

A further simplification is possible if we think of subject matters as given by so-called indirect questions: *who owns the most toys*, *how to write a philosophy*

paper, what I did on my 15th birthday, why Napoleon lost at Waterloo, where to get a good cup of coffee. Then to get an intuitive characterization of **m** from the set of propositions that officially constitutes **m**, you can simply ask yourself, to what direct question are these propositions the competing answers? Putting that direct question in indirect form then gives you an intuitive characterization of the subject matter. What for instance is the question completely answered by **there are exactly k stars** for a specific value of k? It's "how many stars are there?" We are dealing then with the subject matter of **how many stars there are**.

Now, Lewis has a theory of subject matters, but that's not to say he has a theory of *sentential* subject matters: of what *S*'s subject matter is for particular sentences *S*. We are told only that **m** is a subject matter of *S* iff *S*'s truth-value "supervenes" on **m** in the sense that each proposition in **m** implies either that *S* is true or that *S* is false. But that gives *S* a huge range of subject matters ranging from, at the low end, **whether S is true**, to, at the high end, the subject matter of **how everything is in every respect**. Thus we are going beyond Lewis in making a proposal about which subject matter **m** deserves to be considered "*the* subject matter of *S*."

What is the proposal? Thinking again of subject matters as given intuitively by questions, it suffices to identify the question to which *S*'s subject matter is the set of complete answers. That question, I propose, is: *why is S true?* To specify *S*'s subject matter is just to enumerate the possible reasons for *S* to be true, in other words to enumerate *S*'s *truthmakers*. There is nothing more to saying what *S* is about than giving the various reasons it might be true. It will also be helpful to have on hand the set of complete answers to *why is S false?* This, in an inspired turn of phrase suggested by John Macfarlane, can be considered, not *S*'s subject matter, but rather its subject *anti-matter*. The two together are *S*'s subject matter in the broad sense. So, *S* and its negation have different subject matters in the narrow sense but the same subject matter in the broad sense since *S*'s subject matter in the narrow sense is $\sim S$'s subject anti-matter and vice versa. The difference mostly won't be important to us, but I'll try to indicate which I mean on the few occasions when it is.

That's one departure from Lewis but it points to another. Lewis in effect assumes that complete answers to a question are always mutually incompatible; otherwise a question's complete answers wouldn't partition logical space. This is a common assumption but it is arguably wrong for certain questions. Who wants gum? Where can one get a good cup of coffee around here? What are two examples of extinct animals? It's in particular questionable for why-questions, for if the explanandum is overdetermined, each of the overdeterminers gives a complete answer.

Joe was carrying a radioactive bazooka stuffed with plastic explosives; why was he kept off the plane? One complete answer is that he was carrying a bazooka. Another is that he was carrying plastic explosives. Another is that he was carrying something radioactive. These answers are obviously consistent. It's no different with the why-questions to which subject matters correspond, ones of the form *why is S true?* The sentence "Snow is white or cold" is true because snow is white – that's one complete answer to *why is it true?* – and also because snow is cold – that's another complete answer. It follows that we should think of subject matters not as sets of incompatible propositions – none consistent with any other – but at most incomparable propositions – none implies any other.

Now we're almost in a position to define content-part. First let me give me the form of the definition. The form is going to be this:

- (3) *B* is part of *A* iff the inference from *A* to *B* is both truth-preserving – *A* implies *B* – and aboutness-preserving – *A*'s subject matter includes *B*'s subject matter and *A*'s subject anti-matter includes *B*'s subject anti-matter.

I say "almost in a position" because there is still the issue of what it takes for one subject matter to include another.

Lewis has a theory of this too. He reasons as follows: it is harder for worlds to agree on a larger subject matter (the 20th century) than a smaller one (the 1920s). Turning this around, a subject matter is larger to the extent that it is harder for worlds to agree in it. explaining agreement on a subject matter as belonging to the same cell, a subject matter is larger to the extent that its cells are smaller; the larger subject matter's cells subdivide the cells of the smaller one. Subject-matter inclusion thus corresponds to the refinement relation on partitions:

- (4_L) Suppose *m* and *n* are lewisian subjects matters. *n* is a lewis-part of *m* iff *m* considered as an equivalence relation refines *n*, or if we consider them as partitions, *n*'s cells are further subdivided by *m*'s cells.

Now subject matters as we're understanding them aren't quite partitions, so Lewis's account doesn't directly apply. But the core of that account and the part I want to retain is that subject matters grow in inverse proportion to their cells; they become larger as their cells shrink, as each of the original cells gives way to a smaller new cell. And that idea continues to make sense even if the cells aren't disjoint and exhaustive:

- (4) n is a part simpliciter of m iff each n -cell--each set of worlds in n --contains at least one m -cell; equivalently each way for worlds to be n -wise is implied a way for them to be m -wise.

What this means for the special case where m and n are sentential subject matters, say of sentences A and B , is

- (5) $sm^+(B) \leq sm^+(A)$ iff each of B 's truthmakers is implied by one of A 's truthmakers.

What it means for the special case where m and n are subject anti-matters is

- (6) $sm^-(B) \leq sm^-(A)$ iff each of B 's falsemakers is implied by one of A 's falsemakers

Let's now use S^t as a variable ranging over S 's truth-makers and S^f as a variable ranging over S 's falsity-makers; that makes everything more sayable and readable. Plugging (5) and (6) into (3) gives

- (7) B is part of A iff
 (i) A implies B
 (ii) each B^t is implied by an A^t
 (iii) each B^f is implied by an A^f ⁵

That completes the definition of part. It tells us, given two hypotheses A and B , whether B is part of A . So for instance *I have a dog* is part of *Everyone has a dog* because any truthmaker for *I have a dog* is implied by one for *Everyone has a dog*, and every falsemaker for *I have a dog* is implied by, is indeed identical to, a falsemaker for *Everyone has a dog*. Compare this with a different case: *I have a dog* vs *Someone has a dog*. This time there is no inclusion because the vast majority of truthmakers for *Someone has a dog* are not implied by any truthmaker for *I have a dog*.

Now we know how to *recognize* content-parts. However it would also be good to know how to *make* content-parts. Why? Our original problem was to say what it is for A to be *true about a particular subject matter*. I assume that A is

⁵ That A implies B permits a further simplification, because a fact implying B 's falsity implies A 's falsity too: presumably then the reason each B^f is implied by an A^f is that each B^f is an A^f . This means we can ignore the words "implied by" in (iii), and read it as saying that each B^f is an A^f . The official definition remains as is but it helps to remember that we can just as well read (iii) as stating an identity. (This relates to the claim earlier that what explains the falsehood of S 's part thereby explains S 's falsehood, where the falsehood of S 's mere consequence is just a symptom.)

true about a particular subject matter iff the part of it that concerns that subject matter is just plain true. But now what is "the part of A that concerns m "?

An example might help. Suppose that we have a strange reading of Genesis which has God tripling the number of stars every day, the process beginning on the third day after creation. The clearest way to express the hypothesis *The # of stars on day $n = (n-3)^3$* . Our reading of Genesis also tells us that God never created numbers, though, so we don't believe in numbers. Really, we say, our hypothesis only partly true, true in what it says about the concrete world. How do we home in on the part of our hypothesis that concerns stars and the like as opposed to the numbers?

How we approach this depends on whether we take ourselves to be looking for a sentence or something non-linguistic that a sentence if there were one would mean. I take it we're looking more for the latter. Falstaff has probably not taken care to keep his statements about Jones separate from his statements about Smith and the rest of the gang; he was talking about them all together and there's no guarantee that a sentence can be constructed that sums up what he said just about Jones. Likewise there's no guarantee that a sentence can be found that sums up what *The # of stars on day $n = (n-3)^3$* says about the stars.

What we are looking for then is a propositional content, whether that content is expressible in English or not. A propositional content for us will be amalgam of a truth-conditional content B and subject matter b : of a lewisian intension, on the one hand, and a number of subsets thereof which explain why the first intension is true or false. The question is, given

One, a sentence A with such and such a propositional content, and

Two, a lewisian subject matter m (a partition of logical space),

can we find

Three, a propositional content $B (= \langle B, b \rangle)$ such that a sentence *if there were one* that expressed B would count as the part of A that concerned m .

The answer is that we can. It's pretty clear how to do it; you form B by making A blind to distinctions between m -equivalent worlds. The recipe is as follows:

Step One: let B be true in all worlds m -equivalent to worlds in which A is true, and false in the rest.

Each of A 's truth-makers A^t is similarly expanded:

Step Two: let B^t be true in all worlds m -equivalent to worlds in which A^t is true, and false in the rest.

As for B 's falsemakers, they're just a subset of the facts making A false:

Step Three: let the B^f 's be those of the A^f 's that imply the falsity of B ; that is, whose members are not m -equivalent to worlds where A is true.

This means that B is false just in case A is false *for reasons that can be fixed only by changing how matters stand wrt m* . The proposition B we're looking for is the truth-conditional content given by (i) seen as directed at the subject-matter given by (ii) and subject-anti-matter given by (iii). Take the case where A is *The # of stars on day $n = (n-3)^3$* , and the subject matter m is **the concrete world**. The part of A that concerns m is

One, true in worlds concretely indistinguishable from worlds where *The # of stars on day $n = (n-3)^3$* ;

Two, true for the reasons *The # of stars on day $n = (n-3)^3$* might be true, once those reasons have been weakened so as not to distinguish concretely indiscernible worlds;

Three, false for the reasons *The # of stars on day $n = (n-3)^3$* might be false, leaving aside reasons that do distinguish concretely indiscernible worlds—for instance, the fact that there are a hundred stars on the fifth day, but not the fact that there fail to be any numbers.⁶

The part A_m of A that concerns m is true in worlds where A is false, if it is, for non-concrete reasons only—reasons that can be fixed while leaving the concrete world unchanged. The part A_m of A that concerns m is true because of a fact obtained from some truthmaker A^t for A by adding in worlds where A^t is false for non-concrete only—reasons that can be fixed while leaving the concrete world unchanged. The part A_m of A that concerns m is false because of a falsemaker for A that obtains, when it does, for concrete reasons—reasons that can't be fixed while leaving the concrete world unchanged.

⁶ Gideon Rosen in his thesis offered a similar-looking recipe for extracting a mixed mathematical statement's nominalistic content. The Rosen recipe is a special case of the recipe for extracting the portion of any statement's content that concerns a given subject matter.

That completes the construction. It shows that partial truth can indeed be understood naively as truth of a part--where parts are, or can be, obtained by restricting the whole to a certain part of its subject matter. The question is: Why bother with partial truth so defined? Why utter false sentences with true bits rather than just the true bits? The answer was given a long time ago by William James (1979, 31-32):

a rule of thinking which would absolutely prevent me from acknowledging certain kinds of truth if those kinds of truth were really there, would be an irrational rule.

This is usually read as a plea for epistemic boldness; if "acknowledging certain truths" carries a risk of acknowledging the odd falsehood, well, that's a price worth paying. But one can also hear it as a plea for semantic boldness. Suppose that certain truths are best or only accessed as scattered parts of larger falsehoods (or larger hypotheses that might be false). (Quine quote about convenient myths containing the truth as a scattered parts) Dallying with the larger falsehoods may be a price worth paying. The difference with James is that it's not the falsity of one statement tolerated for the sake of another's truth, but the falsity of this statement tolerated for the sake of the truth, or anyway partial truth, of the very same statement - for the sake of the truth this statement contains.

Of course, this apology for partial truth is only as good as the premise that certain truths are only or best accessed as part of larger falsehoods. One can certainly see how the premise *might* be true. Look at in the construction of the true part a moment ago; the construction yielded a true *meaning* but not a true *sentence* whose meaning it was. The only *sentence* available is the original one which we're supposing is or may be false. One can describe the intended meaning and endorse it but there is no obvious way to assert it. What option do we have then but to put the possibly untrue sentence forward in a quasi-assertional spirit? Our plea to the charge of speaking falsely will have to be "guilty with an excuse": part of what we said was true, we don't know how to assert just that part, and we did our best to clue you in to which part it was; it's the part about such and such a subject matter.

So we can see how there *might* be a Jamesian justification for speaking "the truth and something but the truth." It would be nice to have some actual examples though. I'm not sure I have any actual examples but here are some possible examples.

Applied math

Suppose I am not a platonist. I do not believe in numbers and functions. Like anyone else I want to use mathematical physics to describe the physical world; I want to be able to say, for instance, that a projectile's escape velocity is such and such a function of how much it weighs. The escape velocity formula is not true in my view, just as Falstaff's testimony was not true. But it is true in what it says about the physical world. I don't know how to express what it says about the physical world otherwise than to say: the part that concerns the physical world. It is true that my construction of that part runs essentially through platonist worlds.

Am I committed then to saying numbers are metaphysically contingent? No. A world for these purposes is just a coherent scenario. At most I am committed to platonism being conceptually possible. And that's a commitment I happily own up to. Platonism may be untrue, but I never said it was conceptually incoherent! So there's my Jamesian justification.

Negative singular existentials.

I believe that, as I'm tempted to put it, Pegasus fails to exist. But I am not a Meinongian: Pegasus doesn't subsist either; it has no kind of being whatsoever; "Pegasus" is an entirely empty name. This puts me in a bit of a bind, since sentences with empty names in them are false, or anyway untrue. If *Pegasus fails to exist* is in my view untrue, why do I say it?

Well, it is *partly* true: it is true about the things that exist. To go by our definitions above, its truth about the things that exist consists in the full truth of that part of *Pegasus doesn't exist* that concerns the existing things. This is obtained in just the way one obtains the part of *The # of planets = 8* that concerns the concrete things. Recall that the part of *The # of planets = 8* that concerns the concrete things is a proposition true in all worlds concretely equivalent to a Platonic world in which the # of planets really is 8. Likewise the part of *Pegasus doesn't exist* that concerns the existing things is a proposition true in all worlds with the same existing things as a Meinongian world where Pegasus really does not exist. That part is true because although it is not true in our world that Pegasus is nonexistent, it is true in a world otherwise like ours but with Pegasus tacked on as a subsisting thing.

The upshot is that when I say *Pegasus fails to exist*, the part of that that I care about, the part that concerns the totality of what exists, is perfectly true. And since I know no other way to put that part forward than to say *Pegasus fails to exist*, that's what I say. It is true that my construction of the true part runs essentially through Meinong worlds, worlds where Pegasus subsists. But that's OK for the same reason running through a Platonic world is OK. Meinong may

have been wrong, but the view is not conceptually incoherent! There's my Jamesian justification.

Standards of precision

I believe that, as I'm tempted to put it, I am 5 feet 9 inches tall. This puts me in a bit of a bind because I think the statement *I am 5 feet 9 inches tall* is false. For I am less than 5'9" tall (closer to 5'8¾") and you can't be 5'9" tall and less than 5'9" tall at the same time. I believe that, as I'm tempted to put it, there were a million men in attendance at the Million Man March. But I also think that "a million" means "exactly a million" – it entails, for instance, "not ten less than a million." So my statement is almost certainly false. How then can I make the statement?

Well, my subject matter when I said I was 5 foot 9 was how tall people are in inches, in the sense of which height in inches their precise height is closest to. My statement was false, but true about the subject matter of height in inches. My subject matter when I said there were a million men in attendance was attendance in hundreds of thousands, or perhaps lower-bounds-on-attendance-in-hundreds-of-thousands. To express the true part directly, I would have to say that 5'9" is the height in inches closest to my precise height. I could, I suppose, but this way of putting it is ugly and inconvenient and require explicitness about something that I may find it difficult to state explicitly. Better to stick with the original precise statement and let the standards of precision rise and fall with the subject matter.

Pure math

I'm a non-platonist, let's say this time a nominalist. I think it's false that *There are primes over 10, none of which are even*. Like anyone else, though, I want to be able to say it. Why? Well, if I'm to stick with the program, it's because the statement has a part that I do believe, a part that is interestingly true in my view, and remains so even if numbers do not exist. The problem is, this time it is harder to see what the true part might be. Doesn't it follow from the denial of numbers that, as Hartry Field once suggested, existential numerical claims are one and all false, and universal numerical claims are one and all true? That would seem to leave little room for *interestingly* true parts to false numerical claims.

But the idea that *Primes over 10 are even* is every bit as true as *Primes over 10 are odd* assumes that we are dealing with enumerative generalizations about whatever numbers there happen to be. And I don't know why we would assume this, any more than we'd assume *Objects with no net forces on them accelerate to the speed of light* is an enumerative generalization about whatever objects there happen to be. *Objects with no forces on them accelerate* sounds false,

even if there are no objects like that, simply because accelerating is physically unlawful behavior. It's exactly the same with *Primes over 10 are even*; it sounds false whether there are primes over 10 or not, simply because this is mathematically unlawful behavior for primes other than 2.

How are we to explain the intuitive falsity of uninstantiated generalizations like these? The obvious explanation is that they have two parts, one a law saying how objects of the relevant sort behave qua objects of that sort, the other (perhaps presupposed rather than asserted) to the effect that objects of the relevant sort exist. *There are infinitely many primes* has as its first part *Numbers are of a type to include infinitely many primes*, and as its second part *...and there are some*. Nominalists when they say *There are infinitely many primes* are putting the first part forward as true but not the second. Alternatively we could say they are putting the whole statement forward as true-about-a-certain-subject-matter, that subject matter being the *sosein* of numbers rather than their *sein*.

Is it only nominalists who talk this way? Consider "Carnap's paradox." It is just plain obvious that there are prime numbers over 10, but it is not at all obvious that there are numbers. How can that be? Surely if we know a stronger statement, then we know its weaker implication? Well, but maybe what we are representing ourselves as knowing in the first case is not the full statement but the *sosein* part: *numbers are of a type to include infinitely many primes*. *There are numbers* is different in that the *sosein* part is transparently silly: *numbers are of a type to include numbers*. That can't be what we're putting forward as true, or claiming to know. So we are thrown back on the other part: *there are numbers*. And whether there are numbers is not something we can claim to know just on the basis of mathematical competence.

Knowledge of Entailments

Suppose that *A* implies *B*. Knowing that *A* holds of a certain subject may suffice for knowing that *B* holds of that same subject matter. But it doesn't in general suffice for knowing that *B* is true of a different subject matter. The solution to Carnap's paradox is that the subject matter changed in transit from *There are primes over ten* to *There are primes*.

This is interesting if true because Carnap's paradox is reminiscent of Descartes's paradox. Why am I so much happier asserting that I am standing in a large room than that I'm not a bodiless brain-in-vat? Could it be that there is a change in subject matter between *A = I am standing in a large room* and *B = I am not a bodiless BIV*? Let me attempt a proof of concept for this idea; please don't hold me to the details.

What does it take to know that S ? That depends on S 's propositional content, that is, its truth-conditional content \mathbf{S} --a regular old set-of-worlds intension--and its subject matter \mathbf{s} --its truthmakers \mathbf{S}^t and falsemakers \mathbf{S}^f , especially the latter. Hypothesis:

I know that S iff:

1. it is true that S (\mathbf{S} contains the actual world)
2. I believe that S ,
3. I would not have believed that S , had S been false in any of the ways envisaged in its subject matter (I do not believe that S in the nearest \mathbf{S}^f -worlds to \underline{w} , for each \mathbf{S}^f in \mathbf{s})

What does this say about closure under entailment? For an entailment to preserve S 's subject matter is, I suggest, for its subject matter to be part of S 's subject matter.⁷ If that's right, then to say knowledge is closed under subject-matter preserving entailment is to say it's closed under part-entailment; for an entailment whose subject matter is part of S 's subject matter is precisely a part of S .

Now the subject matter in normal contexts of "I am standing in a large room" consists of the normal ways for it to be true, which let's not bother about here, and the normal ways for it to be false: I am sitting in a large room, or I am standing in a small room, I am not in a room at all, or etc. Presumably then in a normal context I do know that I am standing in a large room; for it's true, I believe it, and had I been sitting in a large room, or standing in a small room, I would not have believed it.

What about "I am not a bodiless BIV"? This is not something we say in normal contexts so the question can only concern its subject matter in skeptical contexts. The salient ways for it to be true are ways for things to be as they seem: what some call the good cases. The salient ways for it to be false are ways for me to be a bodiless BIV, of which the most prominent is the skeptical scenario in which appearances are unchanged. If this is the subject matter, then I do not know that I am not a bodiless BIV; for I would still believe this even if the skeptical falsemaker just mentioned were to obtain.

So far, so good. The funny thing is that *even in skeptical contexts*, I'm tempted to claim knowledge that I am standing in a large room--more anyway than I'm tempted to claim knowledge that I'm not a bodiless brain in a vat. How can that

⁷ If we required identity of subject matters, then conjuncts would not in the relevant sense entail their conjuncts; and if knowledge is closed under anything, it's closed under conjunct-dropping entailments.

be, given that skeptical falsemakers have now presumably have made their way into the subject matter of “I am standing etc.”? To go by the analogy with Carnap, it’s because even in skeptical contexts, the proposition expressed by “I am standing in a large room” has a *part*--the part that envisages just the earlier-mentioned “normal” falsemakers--that I *do* still know, notwithstanding the recently introduced skeptical possibilities. One can think of it more generally as the part that concerns subject matter **n = how matters stand in detectable respects**.

Hold on, though. If knowledge is closed under entailment barring a change in subject matter, shouldn’t I also know the part of “I am not a bodiless BIV” that concerns **n**? Well, but I do. The part of “I am not a bodiless BIV” that concerns **n** is true in all worlds not detectably different from worlds like this one--not detectably different from worlds where I correctly perceive myself to be standing in a large room; it is in other words true in the worlds where I seem to be standing in a large room. And I do know that I seem to be standing in a large room.

The problem is that to know that is to know a triviality; compare knowing that numbers are of a type to include numbers. Someone who wonders whether I know I am not a bodiless BIV is not wondering whether I know a trivial part of that proposition, no more than someone wondering whether I know there are numbers is wondering whether I know a trivial part of that proposition. They’re wondering whether I know the whole thing. And I don’t.

Summing up: Truth-conditions are one factor in meaning; subject matter is a second and potentially independent factor. Subject matters let us define a part/whole relation on contents and thereby a notion of partial truth. Semantic moralists to the contrary, settling for *A*’s partial truth makes sense when although *A* may not be itself true, it has an importantly true part that is otherwise difficult to access. Among those who should arguably be willing to settle are: applied mathematicians who aren’t platonists; existence-deniers who aren’t meinongians; loose talkers; pure mathematicians who aren’t platonists; and knowledge-attributors who aren’t skeptics.

Appendix on Ravens

I said *All Fs are Gs* has a different subject matter than its contrapositive. This distinction in subject matter is hard to make out if *All Fs are Gs* is read as just a universally quantified material conditional – everything is either not an F or a G – for that is the same formula we would use to express that all non-Gs are non-Fs, and it would seem on the face of it that the same formula should have the same subject matter. But then, we shouldn't read it that way; indeed the universalized material conditional is often objected to precisely on the ground that it makes *All ravens are black* partly about non-ravens like myself and *All non-black things are non-ravens* partly about black things.

What is the right representation for *All ravens are black*, if not *For all x, either x is not F or x is G*? Some would tinker with the quantifier, but subject matter issues arise already with the embedded conditional. *X is black, if it's a raven* is true in the very same worlds as *x isn't a raven, if it isn't black*: the worlds where either x isn't a raven or x is black. But they seem to differ in *why* they are true.

Consider Rudy, a typical black raven. What makes it true that *Rudy is not a raven, if Rudy is not black*? There's a considerable feeling it is *vacuously* true; given that Rudy is NOT non-black, the question of what kind of thing it is supposing it IS non-black doesn't even arise. What makes it true that *Rudy is black, if it's a raven*? This does not seem vacuously true at all. This time the fact that Rudy is black seems highly relevant.

So far we've got just a difference in reasons for truth, not a difference in subject matter. But the one difference would seem to imply the other. One can't have sentence X true because of snow's color rather than its temperature and sentence Y true because of snow's temperature rather than its color unless X is about snow's color and Y is about snow's temperature. Similarly *Rudy is black, if it's a raven* would not be true because of Rudy's color unless it was about Rudy's color. And *Rudy is not a raven, if Rudy is not black* cannot be about Rudy's color, or it would not be vacuously true.

Now, if the two contraposed Rudy-sentences differ in subject matter, then their universal generalizations (*For all x, x is a raven if black* and *For all x, x is non-black if a non-raven*) presumably differ in subject matter too. Still, of what possible relevance could this be to Hempel's paradox? The two generalizations are still going to be true in the same circumstances; and all confirmation cares about, you might think, is whether a hypothesis is true, not what it's about.

But remember, we are interested in subject matters for the light they might shed on *parts*. Suppose that the difference in subject matter meant that *All Fs*

are Gs came out with *its* instances (things of the form *x is G, supposing it is F*) as parts while *All non-Gs are non-Fs* came out with *its* instances (things of the form *x is not G, supposing it is not F*) as parts. And suppose that confirmation in the relevant sense--the sense in which a black raven confirms *All ravens are black* but a silver coin in my pocket does not confirm *All the coins in my pocket are silver*--is a matter of probabilifying the unobserved instances; of making it likelier that the next F, and the one after that, will be G. Then the confirmational properties of *All Fs are Gs* and its contrapositive may be expected to differ. Because it's the probabilities of the *parts* that matters, and these due to the shift in subject matter are not the same.