Knowing About Things

1 Kinds of Knowledge

I should start by indicating where I'm heading since we may not entirely get there. I want ultimately to be talking about the relations that obtain between (seemingly) different sorts of knowledge—-seemingly, because it may turn out to be the attributions that differ, rather than the states attributed. The kinds of knowledge at issue are, first, knowledge-that—knowing that plastic forks aren't expensive, say. Knowledge-that, or propositional knowledge, is of course the philosopher's favorite and it'll be my main focus too.

The second type of knowledge at issue includes, for example, knowing when, knowing where, why, what, which, whether, and how. *Knowing-wh* is the term sometimes used for the obvious reason that "where," "why," and so on all start with "wh," "how" and its variants being the exception. *Interrogative* knowledge is another good term since "where," "which, " and so on are words we used to ask questions (all except "whether"). And indeed phrases like "what plastic forks cost" are called questions (indirect or embedded questions) in the semantics literature.

The third kind of knowledge at issue is knowledge *about*. This is the kind of knowledge we attribute when we say Louise knows about buttons, or knows a lot of buttons (in Yinglish, Louise knows from buttons). It's the kind we attribute when we say Louise knows some particular fact about buttons. "Let me tell you this about that," Louise might say, and presumably we should listen to her only if she *knows* this about that. I will call this *topical* knowledge for lack of a better word.

How are the three kinds of knowledge related? The usual practice in philosophy has been to reduce the second to the first. But Jonathan Schaffer has argued in a recent paper ("Knowing the Answer") that this is a mistake and that it is really interrogative knowledge that is the more basic. Knowing that p is more revealingly characterized as knowing that p as the answer to question Q. Of course propositional knowledge attributions aren't explicitly question-involving so Schaffer has to say that there is an unpronounced slot for questions which is filled in some appropriate way by context. Schaffer is thus defending a form of *contextualism* about "S knows that p" which has it expressing that *S knows that p as the answer to Q*, where Q is some contextually indicated question to which p is a potentially correct answer.

What about topical knowledge? This comes in because the formal whatnots that play the role of questions in Schaffer's theory are pretty much the formal objects that play the role of subject matters in David Lewis's work on aboutness. What counts for Schaffer as knowing the answer to a question can equally be understood as knowing how matters stand with respect to a certain subject matter. If one thinks the objects of knowledge come already fitted out with subject matters—if they are what I call "directed propositions"—then knowing how matters stand with respect to a certain subject matter

might itself be understood as knowing a directed proposition full stop. And so we come full circle: knowing that p is on closer examination knowing that p as the answer to some question is on closer examination knowing that p is how matters stand with respect to a particular subject matter is on closer examination knowing that capital P, where capital P is small p seen as directed at the aforementioned subject matter. Propositional knowledge is a kind of interrogative knowledge that can be understood as a kind of topical knowledge that can be understood once again as propositional knowledge (albeit this time with a richer sort of proposition).

2 Tracking

I mention all this only to put it aside for now. Our topic to begin with is a theory of propositional knowledge proposed by Robert Nozick in 1981: the *tracking* theory. The theory says that S knows that p iff

1. S believes that p
2. p
3. S would not have believed p, had it been false (sensitivity)
4. S would (still) have believed p, had it been true (adherence)

This theory is not widely accepted today but it caused a lot of excitement when it first appeared. For it seemed, and to some extent still seems, to address in a unified way a whole bunch of epistemological puzzles.

Probably the signature accomplishment had to do with skepticism. Why is it so much easier to know that I have a hand than that I am not a handless brain in a vat on Alpha Centauri? Well, I would have noticed had I not had a hand. Whereas I wouldn't have noticed were I a handless BIV on Alpha Centauri. My hand-belief is sensitive to the fact it concerns, while my not-a-BIV is insensitive to the fact it concerns.

A second accomplishment, arguably, is how the theory deals with Gettier cases. Suppose Jones and I are up for the same job and I have excellent evidence both that Jones has ten coins in his pocket, and that he is going to get the job. What I don't realize is that I am going to get the job, and that I have ten coins in my pocket. Consider the hypothesis that the successful applicant has ten coins in his pocket. I believe it truly and with justification, but I don't know it. I don't know it, according to the tracking theory, because I would still have believed the successful candidate had ten coins in his pocket even if he hadn't, because my belief is based on the contents of Jones's pocket and those would have been the same had my pocket contained fewer or more coins. That's a failure of sensitivity. There is also arguably a violation of the adherence condition. Because my belief is keyed to the wrong pocket, it does not persist into nearby worlds where the right pocket (mine) has ten coins but the wrong pocket has fewer or more.

A third point in the tracking theory's favor is its response to the lottery paradox. I believe that my lottery ticket will lose and it almost certainly will lose. Why don't I then *know* my lottery ticket is a loser? The tracking theory has a straightforward answer to this. My

belief constitutes knowledge only if it is sensitive to whether my ticket wins or not. And it clearly isn't. I would still have believed, on statistical grounds, that my ticket would lose, even had mine been the winning ticket.

3 Closure

If the theory is now rejected, one main reason for this is that knowledge as Nozick defines it is not closed under known entailment. I can know that p, validly infer from it that q, believe q entirely on the basis of that valid inference, and despite all that be ignorant of the fact that q. (That my p-belief is sensitive to the facts does not mean my q-belief is sensitive.) And this goes against the evident fact that valid inference from known truths seems like a process that *preserves* knowledge.

Nozick would reply that this is not a bug but a feature; the whole point, after all, was to explain the greater intuitive difficulty of knowing that I'm not a handless BIV than that I have hands, even though it is clear all around that my having hands entails that the skeptical counterpossibility does not obtain But most philosophers today think that the greater intuitive difficulty here will have to be explained another way, because if validly reasoning from known premises is not a way of extending knowledge then we might as well just give up and go home. (Kripke on the fallacy of logical deduction.)

This is not a merely theoretical problem, or so Kripke argued in Nozick-bashing lectures given in the mid-1980s. He showed how to construct examples where someone knows, in Nozick's sense, that the structure before them is a red barn, but not that it is a barn. Imagine we are in a corner of fake barn country where red barns are almost always real but barns of other colors are often fake. If this red barn here hadn't been a red barn it would have been a blue barn, and we would have noticed that. But if it hadn't been a barn at all, then the closest remaining option is for it to have been a fake barn; and of course we wouldn't have noticed had it been a fake barn. Maybe we can tolerate some apparent violations of closure, but not this one: not knowing it's a red barn while being in the dark about whether it's a barn.

So, where does this leave us? How are we going to hold onto some kind of closure, while respecting the greater difficulty of ruling out skeptical counterpossibilities?

4 Contextualism

One popular strategy here is to say that if we could keep "S knows that p" expressing the same relation in "S knows she's not a BIV" as it did in "S knows she's got a hand," then the former claim would indeed be true, thus preserving closure. But it is difficult to impossible to maintain semantic stability after the mention of skeptical counterpossibilities. The reason it seems false or questionable that I know I'm not a BIV is that knowing anything becomes a much more demanding affair by the time we get to talking about BIV's. A loose analogy might be this: "A billion dollars isn't a lot of

money [said in the context of arguing about bailout proposals], so three hundred and eighteen dollars and forty cents isn't a lot of money." One may well agree about the billion, but feel that three hundred eighteen dollars and forty cents is actually quite a bit. The contextualist would or could say that the mere mention of smaller amounts, especially the forty cents, flips us into a context where the standards for being "a lot of money" are lower.

This connects up with the three kinds of knowledge as follows. To say that "S knows that p" expresses different relations in different contexts is not yet to say anything about how semantically speaking this comes about. Is it that "knows" is itself a context sensitive expression, with a variable in it to control how stringent a requirement it imposes? Or is it that the thing known, the proposition that p, changes from context to context, from an easily known proposition to one that's harder to know? Or is there perhaps a whole other argument place in knowledge claims at the level of logical form? Schaffer as we saw upholds the third sort of view: knowledge-that is to be understood as a three place relation between subjects, propositions, and questions. He doesn't pull this view out of a hat but tries to argue for it as we now discuss.

5 Reductionism

Schaffer's view is in one way traditional and another way new. Philosophers have been trying for a long time to reduce one kind of knowledge to another. The general motive for doing this is theoretical economy. But there's a particular motive too; if "know" meant one thing before "that" and another before "why" then it should sound funny to say that Louise knows that and why plastic forks are so cheap. That it doesn't suggests we need a reduction one way or the other. That's the respect in which Schaffer's view about propositional vs interrogative knowledge is traditional.

The novelty is in his choice of what reduces to what. Interrogative knowledge has usually seen as a kind of propositional knowledge. I know why water is transparent iff I know that p, where p is the correct answer to "why is water transparent?" I know where to get a good cup of coffee iff I know the correct answer to "where can one get a good cup of coffee?" Schaffer objects to this strategy that it cannot deal with the phenomenon of convergent questions.

Questions are convergent if whatever their other differences, the proposition that correctly answers one also correctly answers the other. So, Q might be "is that a goldfinch in the garden, or a bald eagle?" and Q' might be "is that a goldfinch in the garden or a canary?" Given that the answer in both cases is that it's a goldfinch, the standard account predicts that to know whether it's a goldfinch or a canary is no different from knowing whether it's a goldfinch or a bald eagle; both come down to knowing that it's a goldfinch. And that seems just wrong. To know whether this or that, Schaffer suggests, is to know the answer not de re, as it were, but de dicto; it's to know it *as* the answer to the question of whether this or that. What holds for knowing whether holds

more generally, he thinks. Interrogative knowledge is not knowing a proposition that happens to answer the question Q but knowing it *as* the answer to question Q.

Once again, though, the reductionist was in a certain way on the right track. Knowing is knowing, whether it comes before that or why. If interrogative knowledge can't be reduced to propositional, then propositional should be reduced to interrogative.

(1) S knows that p

is thus more fully spelled out as

(2) S knows that p *as the answer to question Q*,

where the third slot is filled by a question that is somehow hanging in the air when the knowledge attribution is made.

6 Questions and Contrasts

Let's now pursue this idea a little further. We know what S is—a subject—and what p is—a proposition—but what exactly is Q? A question certainly, but what is that?

Hamblin in the 1950s proposed that questions can be identified, for semantical purposes anyway, with the set of their possible answers. If we go along with the usual treatment answers as set-of-world propositions, this makes Q a partition of logical space—the set of possible worlds—into a number of subsets q_1, q_2,.., each subset constituting a distinct answer to Q. Schaffer takes this route too. A knowledge attribution is thus, for him, even more fully spelled out as

(3) S knows that p as the answer to ?{q_1, q_2,....}

Now, though, in a move that we'll be returning to, he further analyzes (3) as a kind of *contrastive* knowledge, of knowing that the correct answer holds as opposed to the disjunction of all remaining answers:

(4) S knows that p as opposed to $q_2 \vee q_3 \vee$ .....

Again, $q_2 \vee q_3 \vee$ ..... is the disjunction of Q's *incorrect* answers, that is, all the answers other than p (assumed here to be q_1).

The move from (3) to (4) loses us information; Schaffer is very up front about this. For p might be the correct answer not only to Q but to another question Q' whose incorrect answers are, though quite different from Q's individually, the same as Q's when you lump them all as in (4). Suppose, e.g., that Q is "is that a dollar bill or not?" and Q' is "is that a dollar bill or a well-done counterfeit dollar or some third kind of non-dollar?" One might have thought that it was easier to know that something is a dollar bill as the answer

to "is it a dollar or is it not?" than as the answer to "is it a dollar or an excellent counterfeit or etc.?" But Schaffer has got to say, and does say, that I know the one iff I know the other. His view too thus seems to run into a problem of convergent questions.

7 Contrastive Closure

More on that later, back now to the larger picture. Schaffer is proposing a reduction of propositional knowledge to interrogative knowledge: what we call knowing that p is on a deeper level knowing whether it is p that holds or rather q. This is supposed to help with skepticism in the following way: knowing whether I have a hand or not would seem a whole lot easier than knowing whether I have a hand or a well executed *fake* hand, and that in turn would seem easier than knowing whether I have a hand or that is just the lie I am being fed by the vat I am floating in.

But now, it is one thing to arrange for a split decision on skepticism—Nozick did that already—–another to combine this with a plausible closure principle on knowledge. Schaffer attempts this in a separate paper. He suggests four closure rules of which let me focus on two (the other two are duals of these).

> *Contraction*: if S knows whether p or q, and q is implied by r, then S knows (is in a position to know) whether p or r.
>
> *Intersection*: If S knows whether p or q, and whether p' or q, then S knows whether p & p' or q.

Contraction s says that if I know that p as opposed to a *weaker* alternative, I am in a position to know it as opposed to a *stronger* alternative. Intersection says that if I know that p and that p' as opposed to the same q, then I know their conjunction as opposed to q.

Contraction on the face of it would seem to give the skeptic an opening. For it implies that knowing whether I have a hand or not puts me in a position to know whether I have a hand or only a vat-image of a hand. (r = I have only a vat-hand entails q = I have no hand.) Schaffer would deny us knowledge of hand vs vat-image of hand, so it seems he must deny us knowledge of hand vs no hand.

But, granted that we want to throw the skeptic a bone of some sort—there are some skeptical counterpossibilities we really can't eliminate—–surrendering knowledge of whether I have a hand seems like throwing him more than a bone, more like the keys to kingdom. A bone would be: fine, so I don't know I have hands as opposed to perfect hand-illusions. Perhaps that will give the skeptic enough of what he wants that he allows us our knowledge whether I have a hand or not. Schaffer seems to be looking for a different kind of split decision: we surrender knowledge of whether we have a hand in hopes of protecting our knowledge that we have a hand as opposed to....what? Apparently as opposed to an easily detectable hand-alternative like an ordinary undisguised stump.

8 Elimimating Possibilities

How does Schaffer reach the result that Contraction holds, thereby depriving us of the knowledge of whether we have hands? All of the closure rules are justified in terms of a familiar sort of Venn diagram. The question of whether p or q is represented by the regions outlined in red and in green. The question of whether p or r is represented by the red region again and the region outlined in blue:

.

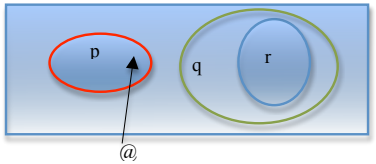

I know that p holds as opposed to q iff I am in a position to eliminate the q-region, where to eliminate a region is to eliminate every world in that region. If I can do that with the q-region, Schaffer reasons, then certainly I can do it with the r-region; indeed I've already done it since every r-word is a q-world. Now we see why I'm in no position to know whether or not I have a hand. That would require the elimination of every single world where I lack one, including the far-away ones that Nozick taught us to think were irrelevant, such as the ones where I'm a brain in a vat on Alpha Centauri.

Now, one question here is whether knowing whether p or q requires me, intuitively speaking, to eliminate every possible q-world. We've been seeing that leads to some surprising results. Another question is, not whether eliminating every possible q-world is necessary for knowing whether p or q, but whether, as Schaffer also assumes, ruling out every possible q-world *suffices* for knowing whether p or q. The idea that it does leads to some even more surprising results, as we see when we look at Schaffer's justification of the Intersection rule. The questions of whether p or q, and of whether p' or q, are represented as follows:

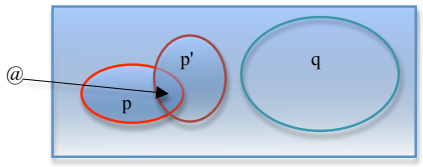

Knowing that p rather than q means that I can eliminate the q-region, and knowing that p' rather than q also means that I can eliminate the q-region. But eliminating the q-region suffices on Schaffer's account for knowing whether p&p' or q, which was the desired result.

What I find puzzling in this justification is that it uses no facts about p and p' beyond that they are disjoint from q. Because it follows from this the very same justification can be given for knowledge of *any* proposition with those properties as opposed to q. The inference rule that is licensed by Schaffer's procedure is therefore a lot stronger than the one he defends with it. The rule is

Explosion: If S knows whether p or q, then S knows whether *anything* or q.

Knowing that this is water as opposed to ink ought to suffice, then, for knowing that this is *so and so many molecules of* water as opposed to ink. It suffices for knowing that this is water and space is curved as opposed to this being ink. Indeed we can leave out the water part; it suffices for knowing that space is curved as opposed to this being ink.

10 Tweaking

Postulating an unarticulated question (or contrast, or something in that vicinity) seems like a good idea in principle. But Schaffer's may not be the best implementation of the idea. A better contextualism, supposing we wanted to go in the same general direction as Schaffer, would be in three ways different from his.

A better contextualism wouldn't lay the whole burden of knowing whether p or q on the subject's being able to eliminate q. Especially if knowing that p is analyzed as knowing whether p or q. For in that case propositional knowledge carries no obligations whatever with respect to the proposition known; the knower's only obligation concerns another proposition with which the proposition known is contrasted.

Laying the whole burden on just one of the two contrasted items is contrary to the spirit of Schaffer's account in any case. He wants knowledge to be deeply contrastive, yet he explains the contrastive notion via a monadic notion of eliminating a hypothesis full stop. Presumably all the problems that arise for monadic knowledge are just going to re-arise for monadic elimination. Indeed they are arguably the very same problems, for if eliminating q is to be a requirement on knowing that p rather than q, then it is hard to see how eliminating q could require more than knowing that q is false. If you are given that p or q, and you know that q is false, that would seem to suffice all by itself for knowing that p as opposed to q. (I can hardly say: not so fast, buster; you may know that q is false, but that doesn't tell you which of p or q is true until you take the extra step of *eliminating* q.)

This reminds us of the second thing that our contextualism better not do: treat eliminating a hypothesis as eliminating every possible way q could obtain. Requiring that makes it

impossible for us to us know whether do or don't have hands. Anyway the requirement is excessive on the face of it. Knowing that q is false does not require the elimination of ever single q-world—that is part of what drives Schaffer to his contrastivism in the first place. And again, if eliminating q is to be harder than knowing that q is false, then there is no need to eliminate q to know whether p or q; knowing that q is false would seem to be quite enough.

A third thing that our contextualism should not do is collapse knowing p as the answer to a *question* into knowing that p as opposed to the *disjunction of the question's other answers*. This is because of the problem of convergent questions we noticed earlier. Knowing whether Alex is a man or a mouse is easy; knowing whether they are a man or a mouse in a brilliant disguise or a mouse not in a brilliant disguise is harder. The disjunctive collapse move obliterates the distinction here because Alex is a mouse in the very same worlds as he is a brilliantly disguised mouse or alternatively a mouse not in brilliant disguise. This is OK with Schaffer because he thinks to know either thing is to eliminate the very same worlds, the ones where Alex is a mouse. But this is an artefact of his theory of elimination which we have already seen reason to question.

11 Interrogative Tracking Theory

So—a better contextualism still in Schaffer's spirit would (a) undo the disiunctive collapse move, and treat propositional knowledge not as contrastive but genuinely interrogative; knowing p is knowing it as the answer to Q. It would (b) would construe the ability to eliminate a possibility as something far less demanding—let's say, for argument's sake, as the fact that I would *notice* if the possibility obtained. It would (c) try to give the proposition supposedly known a role in what it is to know that p; it is not simply a matter of ruling out the other answers.

Here's the proposal, or the kind of proposal, that the above suggests. You can think of it as what Schaffer migt have said if he were more of a tracking theorist, or what Nozick might have said if he were more of a contextualist. I know that p iff

> (1) there is a contextually indicated question Q which p correctly answers,
> (2) if Q is what you're asking, my answer is p
> (3) for each of Q's incorrect answers r, I "would have noticed" had r obtained.
> (4) some sort of positive ground for my belief that p.

This deals pretty nicely with some of the difficulties we've been wrestling with. It is not so hard on this theory to know that I have a hand. For what is the question at issue here? Presumably it's "Do you or do you not have a hand?" Asked whether I did or didn't have a hand, I would answer that I did. This is the correct answer. And I would give this answer in the belief that I would notice had the other, incorrect, answer obtained—had I been without a hand—I would have noticed. But now suppose the question at issue had been "Do you have a hand or do you lack a hand either by having a perfectly undetectable fake hand in the relevant place or in some other way?" (remember the two are

convergent questions for Schaffer). This time I would not be quite so confident in answering that I had a hand. And this is partly because I would by hypothesis NOT have noticed had it been a perfectly undetectable fake hand instead.

12 Interrogative Closure

This is a kind of tracking theory so the question arises what it can say about closure. Let me take a first stab at that. In the literature on questions people sometimes talk about one question being part of another, as "how many stars are there?" seems intuitive to be part of "how many stars are there and how are they distributed?" The definition is

> Q- is <u>contained in</u> Q+ iff each answer to Q+ entails an answer to Q- and each answer to Q- is entailed by an answer to Q+.

I conjecture that closure works something like this:

If S knows the correct answer to a question Q, then S knows (is in a position to know) the correct answer to any question included in Q.

We can try this out on Kripke's closure-type counterexample to Nozick's theory. I can know it's a red barn, Kripke says, without knowing it's a barn, because had it not been a red barn it would have been a different color whereas, had it not been a barn it would have been a perfect papier mache replica.

Here is how things could work so that the one kind of knowledge would suffice for the other; I'll leave you to judge if things plausibly would work that way. The question in the air when we assess me for knowledge that it's a red barn is the question whose answers are (i) yes it is, (ii) nuh uh, it's not a barn, (iii) nuh uh, it's not red. The correct answer is (i). I know (i) as the correct answer only if I would have noticed had one of the other answers obtained. That's certainly true of answer (iii): I would have noticed had it not been red. It's not true of (ii), as Kripke gives the example, since I by hypothesis wouldn't have noticed had it not been a barn. But if we're supposing I know it's a red barn, as we should for this to be a test of closure, then clause (3) of the definition requires us to suppose that I would also have noticed had it not been a barn.

The question at issue when we assess me for knowledge that it's a barn is presumably "is it a barn, or not?" If I would answer that it's a red barn in reply to the previous question, I would presumably answer that it's a barn in reply to this one. Would I have noticed had the other answer been the right one, that is, had it not been a barn? You'll remember that we had to suppose that I would have noticed, to arrange that I knew it was a red barn. So my knowing it's a red barn argues strongly for my knowing that it's a barn. Or to run the example the other way around, the counterfactual *in*sensitivity to barnhood that Kripke postulates would rob me not only of my knowledge that it's a barn but also of my knowledge that it's a red barn.

Of course, we also need to be sure that this style of argument does *not* enable me to draw an antiskeptical conclusion from "ordinary" knowledge like the knowledge that I have a hand. The reason it works in the barn example is this: the one incorrect answer to the question that "it's a barn" addresses is also an incorrect answer to the question that "it's a red barn" addresses. We'll see in a minute that that will generally be the case when one question contains another.

But now, consider the questions "This is a hand" and "This is not a vat-hand" address. They are, let's assume for simplicity, "is this a hand?" and "is it a vat-hand?" The second question is not contained in the first because it has answers — e.g., *it is a vat-hand* — that are not implied by either answer to "is it a hand?," viz. *it is a hand* and *it is not a hand*. So closure as stated above does *not* predict that knowing it's a hand puts me in a position to know it's not a vat-hand. And if you look at the *reason* closure doesn't demand this—the hand question doesn't contain the vat-hand question, on account of the latter having "new" incorrect answers, not implied by any answers to the former—has everything to do with why my knowledge gives out at this point—I would notice if the "old" incorrect answers obtained but this new one would elude me. I won't go into this here, but we can run the reasoning the other way around to show that when the premise's implied question *does* contain the conclusion's implied question, the sensitivity to incorrect answers that gives me knowledge of the premises goes a long way towards guaranteeing me the sensitivity needed to give me knowledge of the conclusion.

13 Topical Tracking Theory

How does topical knowledge—knowledge about—fit in?

Questions are the very same abstract beast as what Lewis calls subject-matters (as Lewis noticed himself). Not only that, but question-inclusion as defined above is very like subject matter inclusion as defined by Lewis. Question inclusion is defined as answers to the first entailing answers to the second, and answers to the second being entailed by answers to the first; subject matter inclusion is the same except that "answer to Q" become "way for things to stand with respect to subject matter **m**. (In both cases we're talking essentially about one partition of logical space refining another.)

This suggests an experiment; suppose we reinterpret the condition that Schaffer calls "S knowing the answer to Q" as being rather a case of S knowing how matters stand with respect to a certain subject matter **m**? Instead of saying that S knows that the barn is red qua answer to "what color is that barn?" (or alternatively qua answer to "what is that red thing?"), we'll now say that S knows that the barn is red as regards the color of the barn (or as regards the identity of the red thing). The toy theory of knowledge now becomes:

> I know that p iff
> (1) p is how matters stand be wrt (contextually indicated subject matter) **m**
> (2) p is how I think matters stand wrt **m**

> (3) I would have noticed had things been different wrt **m**
> (4) some sort of positive ground for my belief that p.

And our closure principle becomes

> If you know p as regards one subject matter, then you are in a position its implication q regarding any included subject matter.

Implications are knowable on the basis of what implies them *barring a change in subject matter*.

To see how this works, suppose you're in fake barn country; it is easy to tell what color a structure is but hard to tell what sort of structure it is that is colored that way. I know that the barn is red as regards the barn's color only if I would noticed had the barn been green or blue. That's easy. I know the barn is red as regards the red thing's structure-type only if I would noticed had the red structure been a church or a papier mache façade. That's hard. Here then is an example where I (plausibly) know a proposition as regards one subject matter but not as regards another. This bears on the closure issue. How can we expect knowledge that p to guarantee knowledge that q despite a change in subject matter when knowledge that p does not even guarantee knowledge that p when the subject matter changes?

14 A Skeptical Puzzle

A more philosophically interesting example is this. Students encountering Descartes's dream argument sometimes remark that their dreams are less vivid and lifelike than waking experience. These students are pretty clear that

> (*) No dream of mine is THIS lifelike.

This may or may not incline the students to reject the dream argument. But although the students may regard (*) as *relevant* to the dream argument, they do not think the argument is thereby *refuted*. They do not claim to know just on the basis of (*) that

> (**) Nothing THIS lifelike—in particular THIS—is a dream.

This is interesting because the hypothesis (*) that they probably know is *logically equivalent* to the hypothesis (**) that they probably don't know. One is of the form *No Ds are L* and the other is of the form *No Ls are D*. Apparently it is easier to know about dreams that that they are not this lifelike than it is to know about experiences this lifelike that they are not dreams, even though a dream is this lifelike only if something this lifelike is a dream.

What is going on here? *No dream is THIS lifelike* is about my dreams, namely d_1-d_{1000}. One of the ways things could have been with respect to them is that d_i (say) could have

been perfectly lifelike. I know that *no dream is THIS lifelike*, then, only if I would have noticed had d_i been perfectly lifelike. A case can be made that I would. More or less lifelike speaks to the felt character of my experience, and that is the kind of thing I am generally pretty sensitive to.

Consider now *Nothing THIS lifelike is a dream*. It's about my phenomenal states, p_1-p_{1000}, individuated by their qualitative character. One of the ways things could have been with respect to them is that p_k (say) could have been a dream. I know that *Nothing THIS lifelike is a dream,* then, only if I would have noticed had p_k been a dream. And I presumably wouldn't, for p_k the dream is subjectively indistinguishable from p_k the waking experience.

15 Reasons to Prefer the Topical Approach (1): Binarity

I've said we *can* repackage the interrogative tracking theory as a topical tracking theory, but of course we can also not repackage it that way. I haven't yet pointed to any advantages of doing so. Let me mention one now before wrapping up.

Knowledge has been binary down through the ages. Schaffer was driven to postulating a third argument place by puzzling shifts in our patterns of attribution; I know I have hands in ordinary contexts, but not when skeptical counterpossibilities are being bandied about.

Now, the topical tracking theory as written also postulates a three place relation, with subject matters taking the place of questions. But the second and third relata—a set-of-worlds proposition and its subject matter—are naturally glommed together into what might be called call a *directed* proposition. A directed proposition is a set of worlds along with a specification of why particular worlds do or don't fall into that set—together, in other words, with a set of truthmakers and falsemakers for the coarse-grained proposition constituted by that set of worlds. Knowledge is widely held to be a relation to fine-grained propositions anyway; directed propositions can be seen as just more flexible and pragmatically sensitive way of going fine-grained. Insofar as "The BARN is red" strikes you as making a different claim than "The barn is RED," it's the directed proposition you're picking up on.

What about the shifts in our patterns of attribution that led Schaffer and others to postulate a third argument? They will be explained somewhat differently on our approach. "I have a hand" still expresses a harder-to-know proposition in skeptical contexts than in ordinary contexts. But the proposition is harder to know, not because the bar has been raised, but because the proposition itself changes. Just as "The barn is red" expresses a harder to know proposition when the topic turns from colors to structures "I have a hand" expresses a harder to know proposition when the topic turns to evil demons and brains in vats.

16 Reasons to Prefer the Topical Approach (2): Closure

The binary approach permits a formulation of closure that focuses entirely on what is known, ignoring any supposed third element in the knowledge relation. First a definition: a directed proposition A is *part* of another one B iff the inference A, therefore B is both (i) truth-preserving and (ii) subject-matter preserving. Closure now becomes:

(I) Knowing a thing puts one in a position to know its parts.

There is still the problem of explaining how knowledge can expand through deduction of consequences other than parts. This is a tricky business but let me say briefly why I don't think the problem is hopeless.

Suppose Louise knows that a certain barn is red. It follows from the barn's being red that the barn is red or leaky. Unfortunately being red or leaky is not part of being red. So closure as just explained does not tell us that knowing the barn is red puts Louise in a position to know the barn is red or leaky.

But perhaps this it not such a bad result. Louise might not realize that the barn is red or leaky if it is red. It's only if she knows the conditional that we would expect knowing the antecedent to give her access to the consequent.

Let's suppose, then, that Louise knows the conditional as well. She knows that the barn is red and she knows that if so it is red or leaky. To be in a position to know on this basis that the barn is red leaky, she has to be able to put these two pieces of knowledge together. That will require another closure principle, this time a multiple-premise principle:

(II) Knowing the parts puts one in a position to know the whole.

Insofar as a conjunction's conjuncts are its parts, it follows by (II) that Louise is in a position to know that the barn is red and if red it is red or leaky.

So far so good, but how does that help with our original problem? After all the conjunction is true in the very same same worlds as *the barn is red* and knowing *that*, we've seen, does not suffice by (I) for knowing that the barn is red or leaky.

But remember—that two sentences are true in the same worlds doesn't mean they express the same directed proposition. The conjunction certainly looks like it has a larger subject matter, involving leakiness as well as redness. Let's say it does and that the larger subject matter includes the subject matter of *The barn is red or leaky.* Then (by our definition of part) for the barn to be red and leaky is *part* of its being red and if so leaky.

So Louise knows something (the conjunction) of which *the barn is red or leaky* is part. Now single-premise closure kicks in to assure us that Louise is in a position to know that the barn is red or leaky.

17 Skepticism

Which leaves an enormous question: why can't I in a similar way parlay knowledge that I have a hand, and that if so I'm not a handless BIV, into knowledge I'm not a handless BIV? The beginnings of an answer were indicated above: it's going to have something to do with *I have a hand* becoming a different and harder to know proposition when the BIV scenario is brought in. But that's for another day.

APPENDIX
in which the topical theory is clarified and a third reason to like it is adduced

The topical theory is a work in progress and that's putting it kindly. What we have so far is that I know that p iff

(1) p is how matters stand be wrt (contextually indicated subject matter) **m**
(2) p is how I think matters stand wrt **m**
(3) I would have noticed had things been different wrt **m**
(4) some sort of positive ground for my belief that p.

This needs work. First, something must be said about what counts as a possible subject matter **m** for p. Second, clause (4) is yet to be filled in. Third, our clarification of what plays the role of p's subject matter will force an adjustment in clause (3).

By definition a subject matter **m** is a partition of (some or all of) logical space—or what comes to the same, an equivalence relation on worlds, where worlds are equivalent iff they are just alike with respect to **m**. The question is, when do worlds count as alike with respect to the subject matter of p?

Certainly if p is true in w1 and false in w2, then w1 and w2 are NOT just alike with respect to p's subject matter. Sameness of truth value is thus a necessary condition on **m**-equivalence. But it is nowhere near sufficient. Imagine that p is a disjunction, say *Smith owns a Ford or Smith has only a bicycle* (f v b). Imagine that p is true in w1 via its first disjunct and true in w2 via its second disjunct. That Smith owns a Ford in w1 and has only a bicycle in w2 clearly makes w1 and w2 different wrt the subject matter of *Smith owns a Ford or a bicycle*. More generally p's subject matter cannot be in exactly the same condition in worlds where p is true for different reasons (or false for different reasons).

I would argue conversely that p's subject matter cannot be in a different condition in worlds where p is true (or false) for the same reason; factors irrelevant to the reasons why p has its truth-value play no part in its subject matter. I conclude that

worlds are alike wrt p's subject matter iff p has the same truth-value in them, for the same reasons.

Thinking of subject matters as partitions and of sets of worlds as propositions, p's subject matter becomes simply *the set of all possible reasons for it to be true or false*—the set of its truthmakers and falsemakers, in one sense of those terms.

Subject matter is contextually variable to the extent that our sense of the factors controlling p's truth–value is influenced by how we conceptualize p. (Do we blame the truth of *the shortest spy is the bank president* on the fact that Hector — the spy in question — runs the bank? Or on the fact that Ortcutt — who runs the bank — is a spy than whom none is shorter? It depends.) But in any particular context, **m** looks like this:

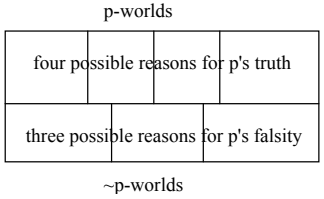

This brings us to the issue about clause (3). Here's what (3) says on the present understanding of **m**: I would have noticed had p been false in any of the ways envisaged by its subject matter (so far so good), AND I would have noticed had p been true in any of the "other" ways envisaged by its subject matter, the ones that don't actually obtain. If by "would have noticed" we mean "would not have believed p," then (3) says that not only would I not have believed p had it been false for any of the indicated reasons, I also would also not have believed it been true for some other reason than the actual one.

And that seems just crazy. To know that there's a duck on the pond, my belief should be sensitive to the various ways there might fail to be a duck, but not to its being *that very duck*, so that I cease to believe it if Daffy is switched for another duck. This suggests we make (3) into

(3') I would not have believed p, had it been false in any of the ways envisaged in its subject matter,

But now what becomes of the business about alternative truth-makers? Is the belief's relation to alternative truth-makers just irrelevant?

Granted that truth-makers have no role to play in clause (3), which speaks to when the belief would go away, perhaps we can find a place for them in clause (4), which speaks

to why the belief is there in the first place. After all it's a common feature of Gettier cases that the belief is there for the wrong reason; the reason p is believed comes apart from the reason why p is true. There may be more than one option here, but if we're aiming for symmetry we could try

> (4') I would still have believed that p, had it not seemed true in any of the "other" ways envisaged in its subject matter.

Take the case where Jones believes someone in his office owns a Ford because he has seen Nogot driving one. Really the Ford is owned by Havit. Jones's belief may well get by condition (3'); if Nogot is borrowing the car from Havit, then had no one in the office owned a Ford, Nogot would not have been seen driving one. (4') puts up more of a fight. That Nogot owns a Ford is one of the "other" (one of the non-obtaining) truthmakers, so (4') requires that Jones would still have believed someone in the office owned a Ford, had Nogot not seemed to own one.[1]

The interrogative theory, as we have seen, lumps all the p-worlds together; it doesn't care about the reasons why p is true, only the reasons why it is false. Insofar as truthmaker structure can be useful in theorizing about knowledge. the topical theory can claim an advantage here. One final example to illustrate this.

Consider p1 = *There is radioactive jade* and p2 = *There is radioactive jadeite or there is radioactive nephrite.* Testing a sample of jade and finding it to be radioactive gives me knowledge that p1. It doesn't matter that my testing device classifies the sample as jadeite when it is really nephrite; it doesn't matter that only nephrite can be radioactive. These added details do, however, seem potentially relevant to my knowledge that p2, for they have the result that I believe p2 "for the wrong reasons," reasons that run through the false disjunct rather than the true one.

The topical theory can account for the epistemic difference here; it is due to the fact that a disjunction's subject matter will find it hard to ignore the difference between truth via the one disjunct and truth via the other. The interrogative theory cannot, because the contrast is all on the truthmaker side and that theory draws no distinctions among a sentence's truthmakers. The contrastive theory has trouble with convergent questions; the interrogative theory has trouble with convergent subject matters.

[1] Obviously this is not perfect. Perhaps I would still have believed that someone owned a Ford if Nogot was not driving one, because in that case Havit would have been. Probably nothing is going to be perfect. But I will leave the matter there.