

Knowledge-wh

Topics in Epistemology
Oxford, 26/4/12

Daniel Rothschild
daniel.rothschild@philosophy.ox.ac.uk

Attributions of knowledge-wh(finite):

- (1) Susan knows who will win.
- (2) Jerry knows which wire controls the device.

Attributions of knowledge-wh(infinitival)

- (3) Mary knows who to call.
- (4) Jerry knows which wire to cut.

...

Here are a couple of questions about the states attributed by these sentences:

- What is their relation to the states attributed by standard attributions of knowledge-that (e.g. *John knows that Bill is at home*)?
- What is their relation to questions (e.g., *Who will win?*, *Which wire controls the device?*).

1 One kind of knowledge

What would be desirable on theoretical grounds, is to have one kind of knowledge, and to account for both knowledge-that and knowledge-wh in terms of it. Moreover, we want this to be compatible with the received view that knowledge consists in a relation between an individual and a proposition. Otherwise, we are in serious trouble: either there is sui generis source of knowledge-wh or all knowledge (that and wh) is not what we thought it was.¹

For this reason REDUCTIONISM about knowledge-wh is an important claim.²

¹Schaffer [2007] holds the latter view, George [2011] argues for something related to the former, as we shall see.

²This is close to what Schaffer [2007] calls the REDUCTIVE VIEW, and what George [2011] calls REDUCIBILITY.

Reductionism We can analyze knowledge-wh in terms of only the knowledge relation between individuals and propositions (as ascribed by such sentences as *John knows that snow is white*).

By ‘only’ I mean that we need not refer to other relations between individuals and propositions or similar entities (e.g. abstract questions) in our analysis.

Let me contrast this with a stronger and (less likely, if not entirely indefensible) thesis, which I will call EQUIVALENCE:

Equivalence For any particular state of knowing-wh (e.g. knowing who came to the party last night), there is a proposition p such that being in that state of knowing-wh just is knowing p .

EQUIVALENCE is stronger than REDUCTIONISM since it requires a specific form of reduction. Both are still vague, but will get further clarified as we consider particular reductions/challenges.

2 Linguistic argument?

Some think that there is a linguistic argument for REDUCTIONISM or something like it. You could read this passage as such an argument (though I don’t know if it was meant that way):

The default assumption is that the ascriptions in (1) [ascriptions of knowledge-wh] are ascriptions of the same kind as ascriptions of the form ‘ X knows that P ’. After all, the same verb “know” occurs in all of these constructions. Furthermore, it is a stable cross-linguistic fact that most of the sentences in (1) are translated with the same verb used in translations of sentences of the form ‘ X knows that P ’. [Stanley, 2011]

This is a good argument that knowledge-wh and knowledge-that are *related*, but does that get us anywhere near REDUCTIONISM. As far as I know (which is not far), it is a stable cross-linguistic fact that one can use the verb *see* in both the following senses:

- (5) I see John.
- (6) I see that you are well.

Clearly, seeing John and seeing that John is here are related. But does one kind of seeing reduce to the other? I doubt it. When one word has two related senses it does not follow that one sense can be reduced to the other in some way (consider also different senses of *city* or *network*). We can conclude, I think, little about reduction from the fact that the same word *know* is robustly used cross-linguistically in different ways. All that *is* (prima facie) shown is a kind of conceptual connection, but not reduction. Another related argument is the argument based on the idea that the same verb ‘know’ is used in both cases of knowledge-that and knowledge-which, and so should be given one meaning. However, giving one meaning to ‘know’ does not establish anything like reduction. Suppose we give the following analysis of the structure of knowledge-wh and knowledge-that constructions.

(7) [John knows [that the blue key works]]

(8) [John knows [which key works]]

We can then think of ‘know’ as a verb which expresses a relation between individuals and certain abstract objects. In the case of that-clauses, it is a relation between persons and propositions and in the case of wh-clauses it is the relation between persons and whatever wh-clauses express (which we will, without prejudice, call QUESTIONS). This surely does not amount to reductionism, it’s just a case in which one word expresses two different relations (or one disjunctive relation).

This is not to say the linguistic facts show nothing. As many emphasize, including Stanley and Williamson [2001], they do push for a syntactic and semantic analysis of knowledge ascriptions of questions in which *know-wh* is not treated as an *idiom*. (Support for this conclusion also comes from other verb-wh constructions such as *wonder-wh*, *tell-wh*, etc.) But the point I have been belaboring is that this alone does not show anything about REDUCTIONISM.

3 A reduction: questions as partitions

Any argument for REDUCTIONISM cannot be based on superficial linguistics facts, we need rather to evaluate a particular reduction. This is our task now.

Just for superficial plausibility, it seems like particular instances might easily be reduced.

(9) *John knows who ran* is true iff for all x such that x ran, John knows that x ran, and for all x such that x did not run, John knows that x did not run.

(9) is a specific instance of reduction of a knowledge-wh attribution. This satisfies REDUCTIONISM since the only relation between people and abstract entities we used was (it seems) the knowledge relation between John and a set of propositions (which we quantify over).

Let us not quibble with this particular analysis, but try to spell out the principles for the general case.

3.1 Questions

We start with an analysis of questions. Here there are a number of classic works by Hamblin, Hintikka, Karttunen, Groenendijk and Stokhof, Lewis, and others.

Rather than presenting these historically, I will simply present an ahistorical amalgam of these to suit our purposes.

By QUESTION I mean not a speech act or a sentence, but a certain abstract object (like a proposition). This is a matter for stipulation, not introspection (nothing feels like question), but once we use the notion, it will be subject to scrutiny.

We associate with any given question a partition of logical space defined by its mutually exclusive complete answers.

Examples:

(10) Who came to the party last night?

Answer 1: Bob came. (not exclusive) Answer 2: Bob and Sam and no one else came. (exclusive)

Note this kind of partition is what Lewis, ‘Relevant Implication’ and elsewhere, calls a subject matter. Questions are subject matters (or topics). This follows normal locution: we talk about questions, questions can be topics of conversation, etc. Extremely important in theory of discourse, certain aspect of focus semantics, etc. We can also define entailment between questions as partitions, etc.

3.2 Knowing-that and knowing-wh

A simple account of knowing-that and knowing-which goes as follows. We assume *know* is a verb that can relate individuals to either questions (i.e. partitions) or propositions. Let q be a partition variable and p be a proposition variable. Let K refer to the knowledge relation between individuals and propositions. Let x be a variable for individuals.

- (11) ‘ x knows p ’ is true iff xKp .
- (12) ‘ x knows q ’ is true iff xKp' , where p' = the cell of the partition q that the actual world is in.

If this analysis is correct we have vindicated REDUCTIONISM.

4 Problems

4.1 Non-exhaustive answers

Groenendijk and Stokhof identified and labelled the mention-some readings of questions:

- (13) Where can you buy cigarettes?
- (14) I know where you can buy cigarettes.

These do provide an argument against the exact reduction given by G+S, but not (seemingly) against reductionism generally. But we could define an ambiguity in question meaning to handle this and maintain REDUCTIONISM. Some semantics needs to be done to handle these, but we’re not doing semantics, so let’s just assume we can do this (which is usually a safe assumption). More on mention-some to come. . .

4.2 Schaffer

Schaffer [2007] argues that reductionism is false. Well, he argues that the REDUCTIVE VIEW is false, and defines that as the view that “knowledge-wh ascriptions should be analyzed via Ksp relations [standard knowledge-relations to propositions] in which the questions Q goes missing.”

His actual examples are hard to apply directly to the view we have above, but similarly structured examples can be made.

Suppose it is common knowledge that John knows that exactly one person was on television last night at 4:45 pm.

Suppose it was Will Ferrell impersonating George Bush. Bill can distinguish Will Ferrell from Janet Jackson but not from George Bush, and just got a momentary glance.

Schaffer claims that these sentences can have different truth values.

- (15) Bill knows whether Will Ferrell or Janet Jackson was on TV.
- (16) Bill knows whether Will Ferrell or George Bush was on TV.

Why is this a problem? On G+S’s semantics these are different attributions of knowledge: to satisfy (15) you need know that WF is on TV and that JJ is not, while to satisfy (16) you need to know that WF is on TV and that GB is not.

The problem comes as follows: we typically assume there is something that is the total knowledge state of Bill, and that this determines which propositions he knows. But then it would seem that either John knows or doesn’t know that Will Ferrell was on TV. If he knows it, then both (15) and (16) are true, if he doesn’t know it they are false. But Schaffer claims that in this case (15) may be true while (16) is false.

Luckily, there’s much to say in response to this. Stanley [2011] argues that this way of reconstructing the problem for G+S’s semantics depends on closure of knowledge (it does) and so is just a way of bringing up typical problems with closure principles. So, people worry about inferring from I know my car is parked outside, to I know it has not been stolen in the last five minutes. Similarly you might worry about inferring (16) from (15).

More can be said than this. (16) and (15) have different presuppositions, and if those presuppositions are imported by the reporters into attributed knowledge the difference in judgment is clear. Another point is there is a complication about the relation between knowing the answer and being able to correctly answer a question. Surely the former does not reduce to the latter (especially as being asked the question could give one knowledge of the Q’s presuppositions by testimony).

I am also skeptical that we can really say:

- (17) John knows whether Will Ferrell or Janet Jackson was on TV but not whether Will Ferrell or George Bush was on TV.

In sum it’s hard to view Schaffer’s example as seriously challenging RE-

DUCTIONISM.

4.3 George

George [2011] makes a more serious challenge to REDUCTIONISM, which will turn out to be important to the question of reductionism about knowledge-how. This relates to the mention-some worry, but significantly extends it. George discusses cases like this:

(18) Benjamin knows where you can buy cigarettes.

George argues that the correct reading for this is that it is true iff Benjamin knows some (nearby) places you can buy cigarettes *and* Benjamin does not falsely believe there are other places you can buy cigarettes.³

I'm inclined to agree with George on these judgments. (You might need to convince yourself of this.)

Given that belief does not reduce to knowledge (even if as Williamson argues—or, at least, asserts—knowledge is more *basic*), then it seems like *some* attributions of knowledge-wh do not reduce to knowledge-that.⁴

Responses will involve delving into the nitty-gritty details of the semantics of questions, unfortunately. I refer you to George [2011], and below, for, what seem to me, effective replies to a variety of strategies for dealing with his case in a way compatible with REDUCTIONISM.

4.4 Identity, related puzzles

Many worries arise when issues of identity emerge. Just as it's hard to know what proposition is ascribed by *Hesperus is phosphorus* it's hard to know what question is meant by *Which celestial body is phosphorus?* As (possibly) with the Schaffer worry it seems like bringing in questions just allows the familiar philosophical worries to arise. See Aloni [2001] for a systematic, descriptivist attempt to deal with all of this.

³Related examples—which do not challenge REDUCTIONISM—are discussed in Klinedinst and Rothschild [2011].

⁴Note that there are no relations to *questions* needed here so this might count as a counterexample to Schaffer's REDUCTIVE VIEW.

5 Equivalence

Recall our stronger thesis, discussed earlier, EQUIVALENCE.

Equivalence For any particular state of knowing-wh, there is a proposition p such that knowing being in that state of knowing-wh just is knowing p .

Well, there is a sort of ambiguity here. Suppose that George (and only George) broke the vase. Then, is knowing that George (and only George) broke the vase equivalent to knowing who broke the vase?

Here's an argument not: it's *possible* to know who broke the vase without knowing that George broke the vase (just think about cases where he didn't!). So the two attributions are not necessarily equivalent. But any two true propositions are extensionally equivalent. So how could EQUIVALENCE possibly be true in an interesting sense? For this reason I am slightly confused when Stanley [2011, p. 37] writes '...it seems that all of the sentences in (1) [ascriptions of knowledge-wh] are synonymous with ascriptions that take the form of 'X knows that p '. Surely the exact point is that when you attribute know which, you don't even need to know what knowledge you are attributing, so *synonymy* cannot obtain, in the usual sense.

Luckily there is a sense in which there's a kind of semantic content of question attributions which *is* equivalent to the semantic content of a propositional knowledge attribution. (Indeed this is somewhat intuitive.) Seeing how this works, however requires using some 2D semantics à la Lewis [1982].⁵

Suppose there's a parameter q which is a world-parameter with respect to the *denotation* of a question is determined. Suppose the denotation of a question is its true (complete, mutually exclusive) answer. The denotation of *Who came?* is, for instance, that Jack and Bill came and no one else did with respect to a q , if Jack and Bill and no one else came in q . Nothing besides questions are sensitive to q .

Now we can call two sentences q -equivalent at w if they are true in all the same worlds if w is set to be the value of the q -parameter.

⁵We could also explore equivalence under the presuppositions of both sentences, Strawson-equivalence, based on von Fintel's [1999] notion of Strawson-equivalence. I think this also is a way to vindicate EQUIVALENCE on G+S's theory or K's.

We can see that EQUIVALENCE (understood in terms of q -equivalence) is true on the account sketched above (the Lewis/G+S account) but would fail for mention-some readings: on the normal understanding of these *John knows where you can buy cigarettes* is true iff there exists a convenient location x such that you can buy cigarettes at x . Clearly, in this case attributing to John knowledge of where you can buy cigarettes is not simply attributing to him knowledge of some particular proposition.

6 Mention-some and domain restriction

We saw that mention-some questions raise worries for both EQUIVALENCE and even REDUCTIONISM (via George).

Can we evade those worries by clever domain-restrictions? After all, it's implausible to think the answer to *Where can you buy cigarettes?* when uttered in Euston Station need make any reference to cigarette-sellers in Oxford. But note the restriction to Euston is not enough. After all I can count as knowing where to buy a cigarette (in the Euston situation) even if I don't know all the shops *in or nearby Euston* that sell cigarettes. (But, as George points out, I need to *not* have false positive beliefs about where cigarettes sold.) Indeed, no single domain restriction would seem to get the truth conditions right for a sentence like *John knows where to buy cigarettes in Euston*.

7 Hamblin/Karttunen

We earlier took the Lewis/G+S view of questions on which they denote equivalence classes dividing up logical space into all the complete answers to a question. This falls in line with Hamblin's [1958] third postulate: 'The possible answers to a question are exhaustive set of mutually exclusive possibilities.'

However, in Hamblin's 1973 semantics for questions—the first—he neglects this and instead defines the meaning of question to be a set of possible, incomplete answers. So for example *Who came?* denotes the set of {the proposition that x walked : x is a person in the domain}.

Karttunen [1977] modifies this slightly. He has a question denote the set of *true* Hamblin answers. Karttunen holds that knowing the answer is knowing all the propositions in the set. So, this is an instance of REDUCTIONISM.

Groenendijk and Stokhof [1984] point out empirical inadequacies. Essentially the worry is that knowing the answer comes too easily. Consider: *Bill knows who came the party*. Does Bill know that simply in virtue of knowing for every attendee that they came? What if he falsely believes that others came? Heim [1994] argues that one should take the Karttunen denotation as primitive and derive the stronger G+S notion. See also Beck and Rullmann [1999] and Klinedinst and Rothschild [2011].

8 Why-questions

I have focused so far on questions where answers are partitioned out in discrete bundles by our naive metaphysics, e.g., *Who came?* Once a domain of individuals (and a contextual meaning for *came*) is fixed the set of answers (in any of the senses above) is clear. However, not all questions are like this. I reserve discussion of *how* till next week, but this still leaves *why*.

The problems are obvious for K, G+S (indeed any theory that depends on having a notion of a complete answer). It is not easy to construct a partition of logical space based on mutually exclusive answers to *why* questions.

Of course, in some cases there will be background assumptions that help us through this, but not always. Richard Feynman's qualifying exam, if I recall correctly, consisted simply of the question *Why is the sky blue?* What partition of logical space would correspond to this question?

Hamblin's second postulate is 'knowing what counts as an answer is equivalent to knowing the question.' Is this plausible?

Analogous, but less severe worries arise with partitioning out time and space for *where* and *when* questions.

9 Reducing question-intension directed states

G+S claim that while verbs like *know* and *tell* (which also embed *that*-clauses) target what they call the question extension (the true answer cell of the partition), other verbs, like *wonder* and *ask*, (which only embed questions) target the question intension (the partition itself).

So while knowing-wh attributes a relation between a person and a proposition, in the end, wondering-wh attributes a relation between a person

and a partition.

This differs from a reductive analysis of asking and wondering in which they are paraphrasable by *want to know*. (I think Hintikka proposes this?) Similarly, we can understand *ask* as paraphrasable as *request to be told*. So perhaps we do not need to posit primitive relations between individuals and questions but can rather build them out of the propositional relations we have.

10 A different reduction

We could try to deal with mention-some worries about reductionism (and the problems with *why*-questions and other questions with less discrete answers) by a less ambitious analysis.

Take something like a classic Hamblin semantics of questions: the meaning of a question is just its set of possible (incomplete) answers. Then we don't build into the meaning of the question any algorithmic specification of a complete answer (and hence don't face the worries about indefinability of *why*-questions, etc).

In order to deal with question embeddings we need to posit a (contextually-determined) answering relation between a proposition and a question, call this relation *A*. pAq iff p is true, p is logically equivalent to the conjunction of some subset of q and p in the context is sufficiently informative to count as an answer to q .

Then our reductive definition of knowledge goes as follows: x knows q iff $\exists p(pAq$ and $xKp)$.⁶ When answerhood requires completeness we get the G+S reading (or the K reading for the weaker notion of completeness) otherwise we can get a mention-some reading.

This is a very loose theory, but one which can accommodate mention-some readings as well as the context-sensitivity of what counts as a complete answer to a 'why' question. It also has little predictive power so needs more stipulations to capture what G+S do, and is more extravagant in requiring two notions: the notion of a question and the notion of being a contextual answer. It's generalizing to the worst case(s)—not my preferred methodology.

What it doesn't give us is an answer to George's worry about reductionism.

⁶I think Higginbotham proposed this as well at some point.

References

- Maria Aloni. *Quantification under Conceptual Covers*. PhD thesis, University of Amsterdam, 2001.
- S. Beck and H. Rullmann. A flexible approach to exhaustivity in questions. *Natural Language Semantics*, 75:249–297, 1999.
- Benjamin R. George. *Question Embedding and the Semantics of Answers*. PhD thesis, UCLA, 2011.
- Jeroen Groenendijk and Martin Stokhof. *Studies in the Semantics of questions and the pragmatics of answers*. PhD thesis, University of Amsterdam, 1984.
- Jeroen Groenendijk and Martin Stokhof. Questions. In J. F. A. K. van Benthem and Alice G. B. ter Meulen, editors, *Handbook of Logic and Language*. Elsevier, 1994.
- Charles Hamblin. Questions. *Australasian Journal of Philosophy*, 36:159–168, 1958.
- Charles Hamblin. Questions in Montague English. *Foundations of Language*, 10:41–53, 1973.
- Irene Heim. Interrogative semantics and Karttunen's semantics for *know*. In Rhonna Buchalla and Anita Mittwoch, editors, *IATLI*. Hebrew University of Jerusalem, 1994.
- James Higginbotham. The semantics of questions. In Shalom Lappin, editor, *The Handbook of Contemporary Semantic Theory*. Blackwell, 1996.
- Lauri Karttunen. Syntax and semantics of questions. *Linguistics and Philosophy*, 1: 3–44, 1977.
- Nathan Klinedinst and Daniel Rothschild. Exhaustivity in questions with non-factives. *Semantics and Pragmatics*, 4(2):1–23, 2011.
- David Lewis. 'Whether' report. In *Philosophical Essays Dedicated to Lennart Åqvist on his Fiftieth Birthday*, pages 194–206. 1982.
- Jonathan Schaffer. Knowing the answer. *Philosophy and Phenomenological Research*, 75:383–403, 2007.
- Jason Stanley. *Know How*. Oxford University Press, 2011.
- Jason Stanley and Timothy Williamson. Knowing how. *Journal of Philosophy*, 98: 411–444, 2001.
- Kai von Fintel. NPI-licensing, Strawson-entailment, and context-dependency. *Journal of Semantics*, 16(1):97–148, 1999.