Exhaustivity and question embedding I

PLING228, 22 January 2016

Goals

Further introduce approaches to question semantics, being to comare them in particular with respect to embedded questions (*know who came*, etc.).

note: issues to be put off til next week and beyond: *which*-questions and de re vs. de dicto readings.

Review

Broad view of question semantics, originating from Hamblin (nb. 1973) through present (we return to Hamblin 1958 below):

- Question semantic values as (characteristic functions of) sets of answers=declarative sentence meanings.
- Declarative semantic values: propositions, modelled as (characteristic functions of) sets of possible worlds=points in logical space/ways things could be

What's an answer? Different views, but commonality is that tied tightly to corresponding declarative sentence (frame).

- Hamblin (1973): "possible answers".
  $[\![$Did John come$]\!]$: $\{[\![$John came$]\!], [\![$John didn't come$]\!]\}$
  $[\![$Who came$]\!]$: $\{[\![\text{x came}]\!]^g : g \text{ an assigment function}\}$
- Karttunen (1977): *true* possible answers (in w).
  i.e. (maximal) subset of above containing only propositions that are true (in w)

Since what possible answers are true differs from world to world, note that the semantic value does too for K (not H).

But we can go from H to K, and other way.

(1) $_H[\![\text{Who came}]\!] = \bigcup\{_K[\![\text{Who came}]\!]^w : w \in W\}$
"Hamblin semantic value = the unioning together of all sets that are K. semantic values in some w"

Questions and answers

- "Pragmatically speaking, a question sets up a choice-situation between a set of propositions, namely, those propositions that count as answers to it." (Hamblin 73)

But what does this mean? As you may have noticed, learning the (true) Hamblin answers = propositions in K-semantic value, does not, intuitively, resolve the question. This is illustrated in (c) below.

Rather, resolving a question seems to equate to knowing what cell we are in in a partition à la Hamblin 1958. This is illustrated in (b) below, which clearly settles the question and corresponds to a cell.

(2) Who passed the exam?
 a. John did.
 b. John and only John did.
 c. John did...
  ...and maybe one other person too.

It might be tempting to think that this argues for a partition semantics. Here is G&S's version of a partition semantics

(3) $_{G\&S}[\![\text{Who came}]\!]^w$
$= \{w' : \forall x(x \text{ came in } w \leftrightarrow x \text{ came in } w')\}$
or: $\{w' : [\![\text{came}]\!]^w = [\![\text{came}]\!]^{w'}\}$

Like K's semantics but unlike Hamblin 1958 this makes the semantic value a (world relative) proposition. But unlike K's this proposition is a cell of Hamblin's partition, i.e. the partition of the set of worlds according to the extension of 'came'. So we can recover the partition in the obvious way.

What does G&S's semantics buy here? According to it, (b) but not (a) is a (semantic) answer, which might seem right.

But it does not explain why (a) appears to resolve the question. (And of course, (b) determines what the true answer is in the H&K sense.)

G&S proposed a "pragmatic" process whereby (a) gets interpreted as (b). Pragmatic explanations have also been proposed within H&K's approach (e.g. Spector 2006) – here the basic idea is that a cooperative, informed speaker would specify *all* true answers so (a) must be the only one.

So data as (1) are likely indecisive.

Notice that K has a natural counterpart that like G&S makes the semantic value a proposition, the conjunction of the true answers; we call this the *K-denotation*:

(4) $_{K'}[\![\text{Who came}]\!]^w$
$= \{w' : \forall x(x \text{ came in } w \rightarrow x \text{ came in } w')\}$
or: $\{w' : [\![\text{came}]\!]^w \subseteq [\![\text{came}]\!]^{w'}\}$

Adopting K', you might think, allows a transparent account of (a), by assuming the following pragmatics: posing a question prompts the hearer for its answer – nb. now unique; a cooperative responder would signal otherwise if he wasn't doing that.

But the situation is more complex in the general case so arguably something more is needed. Even (clearly) partial answers, it turns out, (can) give rise to a kind of exhaustivity inference.

(5) Who came?
  a. At most three boys.
  $\rightsquigarrow$ Nobody else came.

See Spector 2006, van Rooij & Schulz 2006 for attempts to cover both (a) and such exhaustivity inferences in a general way, pragmatically, in the two frameworks.

Existential presuppositions

In asking 'Who laughed at me?' one seems to assume that someone laughed.

Put another way, it can never used to establish whether someone laughed and if so who.

[Similarly: suppose Richard does not know whether Daniel is married, but assumes that if he is it's to Fabienne. He cannot seek to address his ignorance by asking 'Who is Daniel married to?']

This doesn't follow directly on any of the above accounts given that a full resolution on any account requires locating the actual world in the relevant partition (note that one of its cells is the worlds in which nobody laughed).

The are various ways to supplement the accounts by adding a semantic presupposition, or perhaps it's a kind of inference that can be derived given further assumptions (?). I plan to revisit some further relevant data in coming weeks.

Negative responses

...On the other hand, the negative response 'Nobody laughed at you' is licit and obviously resolves the question. This is a closely related issue.

Horn 1972 and Comorowski 1996 analyse this as presupposition denial ('The king of France is not bald').

For G&S it's an answer, so they need say nothing more. In so far as it still counts as one, once the typical existential "presupposition" is analysed (away).

Notice that if the negative response is true, the K semantic value is $\emptyset$ and K' is the tautology, i.e. the set W of all possible worlds. But note that a tautology is not an appropriate response, and in neither case is it obvious why this should be. If it's the strongest (strongest) true answer, asserting it should communicate that nobody came.

The K definitions can be modified bluntly to deal with this (Karttunen himself didn't but did recognise and remedy the problem in another form (see Heim 1994 p.130). Or rely on an existential presupposition (forcing the Horn/Comorowski analysis of negative responses).

Embedded questions

Karttunen was lead to his analysis by the fact that many verbs that embed syntactic questions relate individuals to their answers (unembedded).

(6) a. John knows/predicted/told me/conjectured that Bill and Mary came. ("responsive")

b. John knows/predicted/told me/conjectured-about who came.

(7) John asked/wonders/investigated who came. ("rogative")

Further distinction among responsives: veridical (knows) vs. non-veridical (conjectured-about). Former relate individual to actual true answer, latter to some possible answer. Rogatives, on the other seem to relate them not to any proposition answer but directly to a question. More on all this in coming weeks.

Karttunen focused on veridical responsive verbs and showed they can be analysed, elegantly, in terms of their propositional counterpart

(8) $[\![\text{x knows}_q \text{ Q}]\!] = \{w' : \text{x knows}_p \ {}_{K'}[\![Q]\!]^{w'} \text{ in } w'\}$
i.e. 'x knows Q' iff x knows the proposition which is the K-denotation

G&S criticised this proposal, among other reasons, for failing to explain the following. Suppose John knows that Bill came, but he falsely believes that Mary did. Then (there is at least a sense in which)

(9) John knows who came.

is false. Clearly this is closely related to the issues about exhaustivity above.

G&S proposed their semantics to remedy things. Using ${}_{G\&S}[\![Q]\!]^{w'}$ instead in (8), John has to know that only Bill came which he doesn't.

Heim 1994 re-assessed this argument. She others following her observed that at least with some verbs, there appears to be the reading K' predicts.

(10) a. John was surprised who came (but not who didn't come).
b. John predicted who came.

Suppose John predicts Bill and Mary will come, and says he doesn't have enough info to make further predictions. Still it seems true to say he predicted who came, if only Bill and Mary do. Similar judgment for 'tell' (at least on one reading of 'tell').

Heim observes that in general we can derive the K-denotation from the G&S one, as follows (but not vice versa)

(11) $_{G\&S}[\![\text{Who came}]\!]^{w} =$
$\{w' : {}_{K'}[\![\text{Who came}]\!]^{w'} = {}_{K'}[\![\text{Who came}]\!]^{w}\}$
"the set of worlds in which the K-denotation is what it is w"
= " set of worlds in which the people who came in w came and nobody else did"

Thus there is an argument for using the K-denotation in general, with the particular verb's lexical semantics(s) either using it or using it to derive the K-denotation.

Questions arising: which verbs allow which readings? And why? We will return to this. For now we look at a simpler issue.

Intermediate Exhaustivity

But note that if John in the above scenario had further made some prediction that is false, (10-b) appears to be **false** (or at least has a false reading). The general point is that can be too easy to stand in the V_p relation to the K-denotation so (maybe) Heim's argument wasn't quite right.

What we want is reading the requires John to make no false predictions – the reading we call 'intermediate' exhaustivity because it is weaker than G&S=strong exhaustivity. We propose an operation that can derive the both intermediate and the G&S reading from the basic K-denotation.

Without further constraints then, which there probably are, we might expect a 3-way ambiguity with verbs like predict. (Notice that here a lot of choice points here)

(12) a.
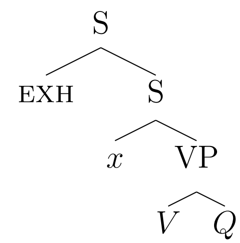

b.
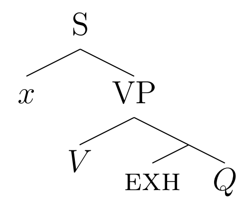

c.
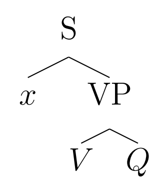

Let us put aside two things til next week. First whether (c) is, after all, a reading for various verbs. Second, whether (a) (intermediate reading) and (c) (G&S) is available for all verbs. As we point out in the paper, (b) cannot be derived without more assumptions for factive verbs like *know* (*surprise*).

(13) a. John predicted that it would rain. $\rightarrow$ it rained
b. John knew that it would rain. $\rightarrow$ it rained

For (12-b) we derive the G&S reading. For (12-a) we derive the following, i.e. intermediate reading:

(14) John predicted K, = the K-denotation of 'who came', but did not predict any possible answer K' to 'who came' that does not entail K
$\approx$ For every x that came, John predicted that x came, and for no x that did not come did John predict that x came

Details

Note: we now make semantic values of questions their intensions directly (i.e. functions from worlds to K-denotation) rather than world relative.

(15) a. $[\![$ who came $]\!] = \lambda w.\lambda w'.\forall x(x \text{ came in } w \rightarrow x \text{ came in } w')$
b. $[\![$ John told predicted Q $]\!] = \lambda w.$ John predicted in w the proposition $[\![Q]\!](w)$

alternatives to questions - $[\![\ .\ ]\!]_F$

(16) $[\![$ who came $]\!]_F = \{Q' : \forall w \exists w' \big(Q'(w) = [\![$ who came $]\!](w')\big)\}$

So we have all the functions which assign to each world a proposition that is the K-denotation of 'who came' in some world, i.e. possible answer but not necessarily the true one.

(17) $P \supseteq Q =_{\text{def}} \begin{cases} Q \rightarrow P, \text{ if } P \text{ and } Q \text{ are of type } t \\ \forall x[P(x) \supseteq Q(x)], \text{ otherwise} \end{cases}$

(18) $P \wedge Q =_{\text{def}} \begin{cases} P\&Q, \text{ if } P \text{ and } Q \text{ are of type } t \\ \lambda x.P(x) \wedge Q(x), \text{ otherwise} \end{cases}$

(19) $\neg P =_{\text{def}} \begin{cases} \sim P, \text{ if } P \text{ is of type } t \\ \lambda x.\neg P(x), \text{ otherwise} \end{cases}$
note: we use $\rightarrow$, &, and $\sim$ to represent the standard truth-function connectives

(20) $[\![$ who came $]\!]_{F+} = \{Q' \in [\![\text{who came}]\!]_F : Q' \supset [\![\text{who came}]\!]\}$

So $[\![$ who came $]\!]_{F+}$ is just the subset of F alternatives to 'who came' that are strictly stronger, in the generalized sense, than the K-denotation. In other words, the elements of $[\![\text{who came}]\!]_F$ that pair each w to a possible answer to 'who came' that is stronger than the K-denotation in w.

[Should have just used $\subseteq$ here!]

exhaustivity operator

(21) $[\![\ \text{EXH}\ \phi\ ]\!] = [\![\phi]\!] \wedge (\wedge_{P \in [\![\phi]\!]_{F+}} \neg P)$
note: where S is a set $\{\phi_1, \ldots \phi_n\}$ of functions of appropriate type, we write $\wedge_{\phi \in S}\ \phi$ for the generalized conjunction $\phi_1 \wedge \cdots \wedge \phi_n$ of its members

strong exhaustivity

(22) a. $[\![\text{EXH [who came}]\!]] = [\![\text{who came}]\!] \wedge (\wedge_{P \in [\![\text{who came}]\!]_{F+}} \neg P)$

b. $= \lambda w.\lambda w'.\forall x(x \text{ came in } w \rightarrow x \text{ came in } w') \wedge (\wedge_{P \in [\![\text{who came}]\!]_{F+}} \neg P(w)(w'))$

(22-b) is the function that maps any w to the proposition that is true in a world w' iff (i) the K-denotation in w holds in w', but (ii) no possible answer to 'who came' holds in w' that is stronger than the K-denotation in w. That is, (22-b) is a function that pairs any w to the set of worlds that agree with w on the extension of 'came'. The latter is just the G&S-denotation of 'who came' in w, and so we derive a strongly exhaustive reading from applying EXH locally to the embedded question

intermediate exhaustivity

(23) $[\![$ John predicted Q $]\!]_F = \{\lambda w.$ John predicted in w the proposition $Q'(w) : Q' \in [\![Q]\!]_F\}$

(24) $[\![$ John predicted Q $]\!]_{F+} = \{\lambda w.$ John predicted in w the proposition $Q'(w) : Q' \in [\![Q]\!]_{F+}\}$
= the set of sets of worlds w such that in w John predicts a strictly stronger possible answer to Q than the one that is actual in w.[1]

(25) $[\![$ EXH [John predicted who came] $]\!]$
$= [\![\text{John predicted who came}]\!] \wedge (\wedge_{P \in [\![\text{John predicted who came}]\!]_{F+}} \neg P)$

Here $\wedge$ and $\neg$ apply directly to propositions (i.e. things of type s,t and so have their customary interpretation. (25) is thus the proposition that is true in a world w iff (i) w is in the denotation of the 'John predicted/told me who came', and (ii) w is *not* in any of the stronger alternatives to this. Condition (i) is satisfied iff John predicted in w the K-denotation in w, as we argued earlier. Condition (ii) is satisfied just in case John did not predict any stronger possible answer to 'who came' than the K-denotation in w.

The two conditions combined then amount to our intermediate reading: (25) is the proposition that John predicted the K-denotation, but did not predict any stronger possible answer

weak exhaustivity

=K-denotation reading. Could be derived, if it exists, by non-application of EXH.

[1] We don't give a general definition of F_+ in the paper. (24) follows if we assume that it is pointwise combination, à la Rooth. Same result appears to follow here if we just assume that it's the stronger (in generalised sense) members of F, since it John's predicting the K-denotation is entailed by him predicting a strong possible answer.