A Uniform Semantics for Embedded Interrogatives: *An* answer, not necessarily *the* answer.*

Benjamin Spector & Paul Égré

(Institut Jean Nicod, CNRS, École Normale Supérieure, EHESS)

Abstract: Our paper addresses the following question: is there a general characterization, for all predicates P that take both declarative and interrogative complements (*responsive predicates* in Lahiri's 2002 typology), of the meaning of the *P-interrogative clause* construction in terms of the meaning of the *P-declarative clause* construction? On our account, if P is a reponsive predicate and Q a question embedded under P, then the meaning of 'P+Q' is, informally, "to be in the relation expressed by P to *some potential complete* answer to Q". We show that this rule allows us to derive veridical and non-veridical readings of embedded questions, depending on whether the embedding verb is veridical, and provide novel empirical evidence supporting the generalization. We then enrich our basic proposal to account for the presuppositions induced by the embedding verbs, as well as for the generation of *weakly exhaustive* readings of embedded questions (in particular after *surprise*).

A semantic theory of interrogative sentences must achieve at least two goals. One goal is to provide an account predicting, for any question, what counts as a felicitous answer, and how an answer is interpreted in the context of a given question. Such an account can then be extended to deal with the dynamics of dialogue. A second goal is to account for the semantic contribution of interrogative clauses when they occur embedded in a declarative sentence. To achieve these two goals, a natural strategy consists in developing a compositional semantics for interrogatives which associates to any interrogative sentence a well-defined semantic value, on the basis of which both the question-answer relation and the truth-conditional contribution of interrogatives in declarative sentences can be characterized in a natural way.

A number of influential theories have such a structure, viz. theories based on a Hamblin-Karttunen type semantics (Hamblin 1973, Karttunen 1977), those based on partition semantics (Groenendijk and Stokhof 1984), as well as theories that unify both approaches (Heim 1994, Beck & Rullman 1999, Guerzoni & Sharvit 2007). In this paper, we are concerned with the second goal, that is we aim to contribute to a general theory of the way embedded interrogatives are interpreted. We will specifically address only a subpart of the relevant empirical domain, namely cases where an interrogative sentence is embedded under a verb or a predicate which can also embed a declarative sentence. As we will see, with one recent exception,[1] no currently available theory is able to achieve what would seem like the main desideratum, namely to derive in full generality, *without specific, idiosyncratic lexical stipulations*, the meaning resulting from embedding an interrogative clause under a predicate that can also embed a declarative sentence.

---

*Acknowledgments to be added.

[1] The exception is Ben George (2011), which partly builds on the ideas we present here. Those ideas were first presented at the MIT LingLunch in 2007 and at the *Journées de Sémantique et de Modélisation* 2008 and in a detailed hand-out form (http://lumiere.ens.fr/~bspector/Webpage/handout_mit_Egre&SpectorFinal.pdf). Our discussions in section V are related to George's work, which we do not discuss extensively.

This paper thus addresses the following question: is there a general characterization, for all predicates P that take both declarative and interrogative complements (*responsive predicates* in Lahiri's (2002) typology), of the meaning of the *P-interrogative clause* construction in terms of the meaning of the *P-declarative clause* construction? The current state of the literature on embedded questions does not provide a systematic answer to this question. Our account of this problem is based on the following idea: if P is a reponsive predicate and Q is a question embedded under P, then the meaning of 'P+Q' is, informally, "to be in the relation expressed by P to some *potential* complete answer to Q". Such a proposal was briefly considered by Higginbotham (1996), but rejected for reasons put forward by Karttunen (1977), which are discussed below. Basically, this gives a weaker semantics to embedded questions than some alternative theories on which the meaning of the embedded question refers to the *actual* answer to Q. The main benefit, we shall argue, is that it allows us to give a principled account of the way in which the lexical semantics of embedding verbs constrains the interpretation of embedded questions.

One clear prediction of such an approach, irrespective of how exactly it is implemented, is the following: a responsive predicate is *veridical* with respect to its interrogative complement (like *know* + question = knowing the true answer to the question) if and only if it is veridical with respect to its declarative complements as well (*know* + declarative entails – in fact presupposes – that the declarative is true). After having reviewed, in section I, the typology of question-embedding predicates, we defend this generalization against apparent counterexamples in section II. These apparent counterexamples involve a class of verbs that we call 'communication verbs', such as *tell*, *announce, predict*,... These verbs appear to license only veridical readings of embedded questions, but are generally considered not to be veridical with that-clauses. We present novel empirical evidence suggesting that these verbs license both veridical and non-veridical interpretations, for both types of complement (declarative and interrogative clauses). In section III, we discuss the notion of *potential complete answer* that we need to use in order to for our proposal to be viable, and we temporarily conclude that we need to refer to a strong notion of complete answer, namely the one that derives from partition-semantics (Groenendijk and Stokhof 1982). At this point in the paper, we thus only predict so-called *strongly exhaustive readings*, as opposed to *weakly exhaustive readings*. Sections IV and V are devoted to two refinements of the basic framework. In section IV, we show how to incorporate certain facts about the presuppositional behavior of certain verbs (like *know*, and *agree*) when they embed questions. In section V, we examine closely the behavior of the verb *surprise*, which has been argued, correctly in our view, to only license *weakly exhaustive readings* (Heim 1994, Beck and Rullmann 1999, Sharvit 2002, Guerzoni & Sharvit 2007, a.o.). We show that by taking into account the *presuppositions* induced by *surprise* when it embeds an interrogative, we can give a plausible account of weakly exhaustive readings. We further show that the move we need to make in order to predict weakly exhaustive readings happens to make fine-grained predictions for *surprise* which turn out to be correct.

I. Veridical vs. Non-Veridical Responsive Predicates

Let us start with Lahiri's typology of interrogative embedding predicates. The following tableau is adapted from Lahiri (2002: 286-287):[2]

[2] The case of "decide" is not in Lahiri's original list. In Egré (2008), *decide* is treated as a rogative non-veridical predicate.

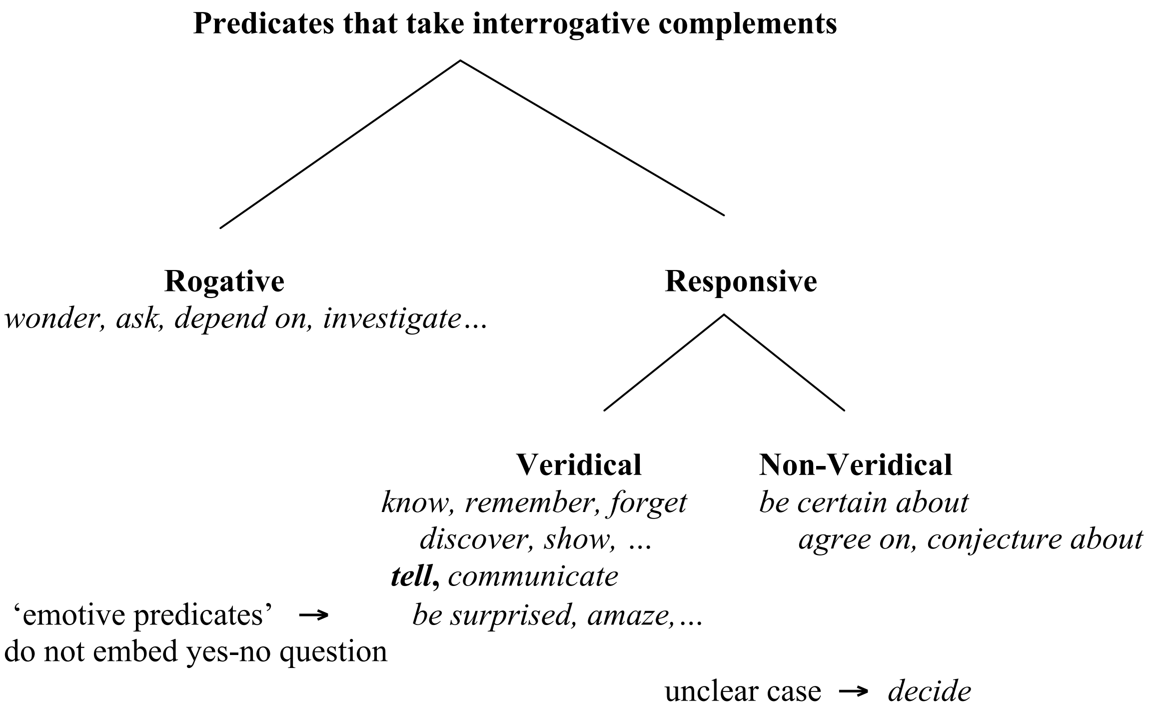

Rogative predicates are predicates that take interrogative complements and do not license declarative complements. We will not be concerned with such predicates.

Responsive predicates, which are the focus of our study, are characterized, among others, by the two following properties:

(i) they take both declarative and interrogative complements
(ii) they express a relation between the holder of an attitude and a *proposition* which is an answer to the embedded question

These two properties are illustrated below:

(i) they take both declarative and interrogative complements

(1) a. Jack knows that it is raining.
b. Jack and Sue agree that it is raining (reciprocal reading: Jack agrees with Sue that…)

(2) Jack knows whether it is raining.
→ Jack knows that S, where S is the true answer to *is it raining?*

(3) Sue and Peter agree on whether it is raining
→ Either both believe that it is raining, or both believe that it is not raining
i.e. : Sue and Peter agree that S, where S is a potential answer to *is it raining?*

(ii) they express a relation between the holder of an attitude and a *proposition* which is an answer to the embedded question

(4) "John knows whether it is raining" is true iff John knows p, where p is the correct answer to "is it raining ?"

(5) "Sue and Peter agree on whether it is raining" is true iff Sue and Peter agree that p, where p is a *possible answer* to "is it raining?"

Lahiri divides responsive predicates into two classes:

veridical-responsive: Responsive predicates that express a relation to *the actual true answer* (= *extensional* predicates in Groenedijk & Stokhof's sense)

non-veridical-responsive: Responsive predicates that express a relation to *a potential answer* (not necessarily the true answer)

Illustration:

(6) Jack knows whether it is raining
→ entails that Jack has a *true belief* as to whether it is raining

(7) Sue and Peter agree on whether it is raining
→ true even if Sue and Peter both believe that it is raining while in fact it isn't

As Lahiri (2002) notes, following Berman (1991), veridical-responsive predicates express a relation between an attitude holder and a proposition that is, in some sense to be made precise, *the actual complete answer* to the embedded question, while non-veridical responsive predicates express a relation between an attitude holder and a proposition that is simply *a potential complete answer* to the embedded question. The first class thus coincides with the class of *extensional embedding predicates* as defined by Groenendijk & Stokhof (1982). While there have been several proposals (Karttunen 1977, Groenendijk & Stokhof 1982, Heim 1994, a.o.) to relate the two properties mentioned above in the case of veridical-responsive predicates, there hasn't been any convincing attempt to explain systematically which responsive predicates are veridical and which are not. The following two questions are thus in order:

(8) a. Can we predict which responsive predicates are veridical and which are not, on the basis of their meaning when they take a declarative complement?

b. Is there a uniform characterization, for all responsive predicates, of the semantic relation between the interrogative complement variant and the declarative complement variant?

By looking at Lahiri's typology, one is struck by the following fact: except for a narrow class of exceptions, all veridical-responsive predicates are factive, hence veridical,[3] with respect to

---

[3]Let's make our terminology clear (though still informal):

- a predicate P is *veridical-responsive* if it can take an interrogative clause Q as one of its argument and is such that [X P Q] (where X is an other argument of P, if P requires one, the null string otherwise) is true iff and only if [X P S] is true, where S expresses the actual complete answer to Q
- a predicate P is *factive* with respect to its declarative complements if it can take a declarative clause S as one of its arguments and is such that [X P S] *presupposes* that S is true
- a predicate P is *veridical* with respect to its declarative complements if it can take a declarative clause S as one of its arguments and is such that [X P S] *entails* that S is true

Assuming that any presupposition of a sentence is also an entailment of this sentence, it follows that a predicate that is factive with respect to its declarative complement is always also veridical with respect to its declarative complement. Note that the generalization that all *veridical-responsive* predicates are also veridical with respect to their declarative complements is by no means a logical necessity (it has actually been explicitly denied, as we will see) their declarative complements, while all non-veridical responsive predicates are neither factive nor veridical with respect to their declarative complements. The narrow class of exceptions involves the verbs *tell*, *announce, predict…* As observed by many authors (Karttunen 1977, Groenendijk & Stokhof 1982, Lewis 1982, Berman 1991, Higginbotham 1996, Lahiri 2002), while (9) below does not entail that what Jack said to Mary is the truth (i.e. does not entail that Peter is the culprit), (10) does intuitively suggest that Jack told Mary the truth as to who the culprit is. Thus *tell* appears to be both *veridical-responsive* and *non-veridical* with respect to its declarative complements. Similar observations hold for *announce* and *predict*.

(9) Jack told Mary that Peter is the culprit.
(10) Jack told Mary who the culprit is.

One feature that these verbs share is that their semantics involves, intuitively, some reference to *speech acts* (that is, for Jack to tell someone whether/that is raining, Jack must have said something; likewise for *predict* and *announce*). We call such verbs *communications verbs*. Mostly on the ground of these exceptions, it is widely assumed that whether a responsive predicate is veridical-responsive or not is to be encoded in its lexical semantics on a case by case basis, independently of whether it is veridical with respect to its declarative complements (see for instance Groenendijk & Stokhof 1993, Sharvit 2002). However, the fact that these exceptions always involve communication verbs suggests that we should look for a principled explanation of their behavior, instead of resorting to mere lexical stipulations.

Putting aside these exceptions for a moment, we would like to defend the following claim, which is an answer to (8)a:

(11) Veridical-responsive predicates are exactly those responsive predicates that are factive or veridical with respect to their declarative complements.[4]

In section III, we will show that this generalization follows directly from a uniform characterization of the semantic relation between the interrogative complement variant and the declarative complement variant, i.e. from an answer to (8)b. Before offering such an answer, though, we want to show that our generalization in (11) can be defended against the apparent counterexamples mentioned above, i.e. the behavior of communication verbs.

II. Communication verbs: new data

The goal of this section is to have a closer look at question-embedding communication verbs. We will argue that a) contrary to previous claims, they are not systematically veridical when they take an interrogative complement, though they tend to be so, and b) that they actually display some kind of an ambiguity when they take a declarative clause as their complement; namely, they can in fact have a factive reading. Taken together, these two facts lead us to conclude that communication verbs do not after all falsify (11).

---

[4] It will turn out that all such predicates are actually factive, and not merely veridical, with respect to their declarative complements, a fact that will become crucial in section V.

II.1 Communication verbs are not systematically veridical-responsive

As observed above, a sentence like (12) usually triggers the inference that what Jack told Mary is the true answer to the question "Who is the culprit?", and thus supports the conclusion that *tell* is veridical-responsive:

(12) Jack told Mary who the culprit is.

Yet this conclusion did not seem entirely obvious, for instance, to David Lewis, as the following passage illustrates (Lewis 1982, p. 46, on the idea that "tell whether p" should mean "if p, tell p, and if not p, tell not p"):

(13) "This is a *veridical* sense of telling whether, in which telling falsely whether does not count as telling whether at all, but only as purporting to tell whether. This veridical sense may or may not be the only sense of '*tell whether'*; it seems at least the most natural sense."

While arguing *for* the view that there is 'a *veridical* sense of telling whether' which is its "most natural sense", Lewis is cautious not to exclude the possibility of a non-veridical reading.

And indeed we find examples in which *tell+question* is not as clearly veridical as, say, *know+question*. Consider the following contrasts:

(14) Every day, the meteorologists tell the population where it will rain the following day, but they are often wrong.
(15) # Every day, the meteorologists know where it will rain the following day, but they are often wrong.

(16) I heard that Jack told you which students passed; but I don't think he got it right.
(17) #I heard that Jack knows which students passed; but I don't think he got it right.

These contrasts suggest that while there is a *tendency* to infer, from 'X told Y Q', that X told Y the actual complete answer to Q, this inference can vanish in ways that are not attested for 'X knows Q'. Similar observations hold for other communication verbs:

(18) Jack predicted/announced whether it would rain.
(19) Every day, the meteorologists predict/announce whether it will rain the following day, but they are often wrong.
(20) I heard that Jack predicted/announced which students would pass, but I don't think he got it right.

While (18) strongly suggests that Jack made a correct prediction, (19) and (20) are felicitous, but should not be so if *predict* and *announce* were systematically responsive-veridical.

II.2. Communication verbs can be factive with respect to their declarative complements

II.2.1. *Tell*

Of course, assuming that *tell*, *predict* and *announce* are not truly responsive-veridical, or that they are ambiguous between a responsive-veridical variant and a non-veridical variant, one would like to know why they tend to favor a veridical interpretation in simple cases like (12)and (18). Even though we are not going to provide a full answer to this question, we now show that communication verbs actually have also *factive uses* when they embed declarative clauses (Philippe Schlenker should be credited for most of these observations – see Schlenker 2006).

Consider the following sentences:

(21) Sue told someone that she is pregnant.
(22) Sue didn't tell anyone that she is pregnant.
(23) Did Sue tell anyone that she is pregnant?

From (21), on easily infers that Sue is in fact pregnant. This is of course no argument for the view that (21) is a case of a factive or veridical use of *tell*, since one can resort to a very simple pragmatic explanation to account for this inference. Namely, assuming that people, more often than not, believe what they say, one infers from (21) that Sue believes that she is pregnant. Given that Sue's belief is unlikely to be wrong in this case, one infers that Sue is pregnant. However, (22) and (23) also yield the same inference. Yet, in these two cases, there is no obvious pragmatic explanation for why this should be so. Indeed, the fact that Sue did *not* say to anybody "I am pregnant" is no evidence whatsoever that she is in fact pregnant. If anything, one could in principle conclude, on the contrary, that Mary is not pregnant, since this could be actually a good explanation for why she didn't say to anybody "I am pregnant". In fact, if Sue were pregnant, chances are that she would have told at least one person that she is. Likewise, the mere fact of asking whether Sue said to anybody "I am pregnant" should not yield the inference that the speaker in fact believes that Sue is pregnant; in principle, he could ask whether she said so in order to know whether she in fact is.[5]

So we have exhibited a case where a sentence of the form *X told Y that S* triggers the inference that S is true and this inference is preserved under negation and question-formation – a "projection pattern" typical of presuppositions. In fact, even in cases where there is no specific pragmatic reason to infer from X saying that S that S is in fact true, we find that the same projection pattern is still possible. Thus consider the following sentences:

(24) Sue told Jack that Fred is the culprit.
(25) Sue didn't tell Jack that Fred is the culprit.
(26) Did Sue tell Jack that Fred is the culprit?

While (24) may or may not trigger, depending on context, the inference that Fred is the culprit, (25) and (26) both strongly suggest, out of the blue, that Fred is in fact the culprit (even though this inference is not present in certain contexts). That we are dealing here with

---

[5] Strikingly, the question « Did Sue say that she is pregnant? does not seem to suggest that Sue is in fact pregnant, at least not to the same extent as (23). According to some informants if a dative argument is added for the verb 'say' (« Did Sue say to anyone that she is pregnant? »), then the question is more easily understood as implying that Sue is pregnant, but sill not to the same extent as (23).

some kind of presupposition is shown by the fact that such sentences pass the *Wait a Minute Test* (Von Fintel 2004), which has been argued to be a good test for presuppositions, and which we now describe.

(27) The Wait a Minute Test ('WMT test' for short)
If S presupposes p, then the following dialogue is felicitous:
- S
- Hey wait a minute! I didn't know that p!
Illustration:
- The king of Syldavia is (not) bald
- Hey wait a minute! I didn't know there was a king of Syldavia

Let us now apply the WMT test to (24), (25) and (26):

(28) - Sue told Jack that Fred is the culprit
- Hey wait a minute! I didn't know that Fred is the culprit

(29) - Sue didn't tell Jack that Fred is the culprit
- Hey wait a minute! I didn't know that Fred is the culprit

(30) - Did Sue tell Jack that Fred is the culprit?
- Hey wait a minute! I didn't know that Fred is the culprit

Interestingly, *say that p* does not behave the way *tell someone that p* does:

(31) Sue said that Fred is the culprit
(32) Sue didn't say that Fred is the culprit
(33) Did Sue say that Fred is the culprit?

In the absence of specific contextual factors, none of the above sentences yields the inference that Fred is in fact the culprit. And in a context in which we take Sue to be well informed, (31) suggest that Fred is in fact the culprit, but (32), if anything, could lead the hearer to conclude that Fred is *not* the culprit (since otherwise Sue would maybe have said so). Some informants notice that stressing *say* may actually make it behave like *tell*.

Whatever the source of this complex behavior is, we may safely conclude that *tell someone that p* has a *factive use*, though it is not *always* factive (as opposed, say, to *know*). We may capture this fact either in terms of a lexical ambiguity or as the by-product of some as yet unknown principle that governs the *generation of presuppositions* (Schlenker, 2010).

Taking into account the facts pointed out in section 2.1, we have shown that *tell* has the following properties:

(34) Properties of *tell*
a) *tell whether* tends to be veridical-responsive, but is not always (contrary to *know whether*, which is always veridical)
b) *tell that* has a factive as well as a non-factive use

These observations are sufficient to maintain our generalization (11), repeated as (35):

(35) Veridical-responsive predicates are exactly those responsive predicates that are factive or veridical with respect to their declarative complements.

Namely, *tell* is no counterexample to the generalization, since it is consistent with all the available data to claim that *tell* is actually factive under one of its readings, and that the factive variant is in fact the one that most easily embeds questions in English (though the non-factive variant may well be able to embed questions, giving rise to non-veridical readings of the type illustrated in section II.1).

For sure, many things remain to be explained, in particular the fact that the veridical-responsive use is clearly the preferred one. Even though we don't have an account of this latter fact, we would like to show that the cluster of properties that we have identified for *tell* carries over to other communication verbs. A more in-depth study is therefore needed for this class of verbs.

II.3. Other communication verbs: *Announce*, *Predict*

Consider now:

(36) Mary announced that she is pregnant.
(37) Mary didn't announce that she is pregnant.
(38) Did Mary announce that she is pregnant?

(36) tends to trigger the inference that Mary is pregnant, which in itself is not surprising (assuming that Mary announces only what she believes is true, and that she knows whether she is pregnant). But the fact that the very same inference is licensed by (37) and (38) is not expected. Again, this is the projection pattern of presuppositions. That we are dealing here with a presupposition is confirmed by a) the projection pattern found in quantificational contexts, and b) the WMT test:

(39) None of these ten girls announced that she is pregnant.

There seems to be a reading for (39) which licenses the inference that all of these ten girls are pregnant, much like "None of these ten girls knows that she is pregnant".

(40) -Sue didn't announce that she is pregnant.
- Hey wait a minute! I didn't know that Sue is pregnant!

(41) - Did Sue announce that she is pregnant?
- Hey wait a minute! I didn't know that Sue is pregnant!

The very same pattern can be replicated with *predict*. Thus consider:

(42) Mary predicted that she would be pregnant.
(43) Mary didn't predict that she would be pregnant.
(44) Did Mary predict that she would be pregnant?
(45) None of these ten girls had predicted that she would be pregnant.

While (42), depending on context, may or may not trigger the inference that Mary got pregnant, this inference is easily drawn, out of the blue, from (43) and (44). And (45) strongly suggests that the speaker takes the ten girls he is referring to to be actually pregnant (or to have been).

Again, the WMT test confirms that we are dealing here with presupposition-like inferences: both (43) and (44) licenses an objection of the following form:

(46) Hey, wait a minute! I didn't know that Mary got pregnant !

These observations, together with what has been established in the previous sections, argue for the following generalization:

(47) Properties of question-embedding communication verbs
Let V be a question-embedding communication verb.
a) *V + question* tends to be veridical-responsive, but is not always (contrary to *know whether*, which is always veridical).
b) *V + declarative* has a factive as well as a non-factive use.

II. 4. *Predict* and *foretell* in French

In support of the generalization stated in (35), we should point out an interesting French minimal pair, made up of the two verbs *predire* ('predict') and *deviner* (approximately 'foretell'). These verbs are close in meaning, with one major difference, namely the fact that *deviner* is factive with respect to its complement clause, while *prédire* is not. There are other differences, in particular in terms of their selectional restrictions,[6] but when they take an animate subject and a complement clause, the resulting sentences assert that the denotation of the subject has made a prediction whose content is that of the complement clause. Another noteworthy difference is that *prédire* tends to imply the existence of a speech act, while *deviner* does not.[7] Now, it turns out that both *prédire* and *deviner* can take interrogative complements. What (35) implies is that *deviner*, but not *prédire*, is veridical-responsive. This prediction is clearly borne out :

(48) (?) Marie a prédit qui viendrait à la fête, mais elle s'est trompée.
Marie predicted who would attend the party, *but she got it wrong.*
(49) # Marie a deviné qui viendrait à la fête, mais elle s'est trompée.
Marie foretold who would attend the party, but she got it wrong.

Although (48) is, for some speakers, slightly deviant out of the blue, suggesting that these speakers have a *preference* for a veridical-responsive use (cf. our discussion of communication verbs, to which *prédire* belongs), there is nevertheless, even for such speakers, an extremely sharp contrast between (48) and (49). And the following judgments, which make the same general point point, are accepted by all our informants :

---

[6] While *deviner* has to take an animate, sentient subject, *prédire*, just like English *predict*, can take any subject whose denotation can be conceptualized as carrying some kind of propositional information. Thus, a linguistic theory can *predict* a certain fact, and the French counterpart of the phrase 'This theory' can be the subject of *prédire*, but not of *deviner*.

[7] This is a subtle contrast. There are uses of *prédire* which do not imply a speech act, as when one says that a *theory* predicts something. But it is at least true that, out of the blue, a sentence with *prédire* is understood to imply the presence of a speech act.

(50) *Chacun des enquêteurs a deviné quels suspects seraient condamnés, mais certains se sont trompés.
'Every investigator foretold [factive] which suspects would be condemned, but some of them got it wrong'.

(51) Chacun des enquêteurs a prédit quels suspects seraient condamnés, mais certains se sont trompés.
'Every investigator predicted [made a prediction as to] which suspects would be condemned, but some of them got it wrong'.

II.5. *tell* in Hungarian

So far, we have argued that communication verbs are somehow ambiguous between a factive reading and a non-factive reading when they embed declarative complements, and that they are likewise ambiguous when they embed interrogative clauses, giving rise to veridical readings as well as non-veridical ones. In this section, we show that some data from Hungarian provide additional support for this view. The relevant facts, which were described to us by Marta Abrusan (p.c.), are as follows. The Hungarian counterpart of *tell* comes in different variants; every variant is based on the same root (*mond)*. We will specifically focus on two variants, *mond* and *elmond* (*el* is a perfective particle). When taking a declarative complement, *mond* is non-factive, but *elmond* is factive, as illustrated by the following judgments (M. Abrusan, p.c.).

(52) Péter azt mondta Marinak, hogy az Eiffel-torony össze fog dölni.
Peter that told Mary.DAT that the Eiffel tower PRT will collapse.
'Peter told Mary that the Eiffel tower will collapse'.
→ No inference that the Eiffel tower will in fact collapse. Given background knowledge, one understands that Peter was fooling Mary.

(53) Péter elmondta Marinak, hogy az Eiffel-torony össze fog dölni.
Peter EL 4 .told Mary.DAT that the Eiffel tower PRT will collapse.
'Peter told Mary that the Eiffel tower will collapse'.
→ The Eiffel tower will collapse.

Now, both *mond* and *elmond* embed interrogative complements. As we expect given our generalization in (35), *mond* is not veridical-responsive, but *elmond* is. The relevant judgments are as follows. (54) shows that *elmond* is veridical responsive.

(54) Péter elmondta Marinak, hogy ki fog nyerni.
Peter el.told Mary.DAT, that who will win.INF
'Peter told Mary who will win'.
→ What Peter said is true.

With *mond* can (but does not have to) be 'doubled' by the accusative pronoun *azt* (which is then stressed), which forces a contrastive reading for the declarative clause ('He said *p*, not *q*'). Whether or not *azt* is present, a non-veridical reading is available (and maybe also a veridical one).

(55) Péter (AZT) mondta Marinak, hogy ki fog nyerni.

Peter it-ACC told Mary.DAT, that who will win.INF
→ There is no implication that Peter told the truth.

II.6. Summary

We have seen that across languages, communication verbs display an ambiguity between factive and non-factive uses when they take declarative complements. As expected given generalization (35), these verbs also create an ambiguity between a veridical-responsive and a non-veridical-responsive use when they embed interrogatives. While we will not venture to make any specific hypothesis as to the source of this ambiguity (in particular as to whether this ambiguity has to be thought of as a lexical ambiguity or as being pragmatically driven), we hope to have shown that such verbs do not undermine the generalization proposed in (35). We also exhibited a French minimal pair *(prédire* vs. *deviner*) and data from Hungarian which further support our generalization.

III. Towards a uniform semantic rule for embedded interrogatives

The gist of our proposal can be summed up as follows:

(56) For any responsive predicate P, a sentence of the form X P Q, with X an individual-denoting expression and Q an interrogative clause, is true in a world w if and only if the referent of X is, in w, in the relation denoted by V to **some** proposition A that is a **potential complete answer** to Q, i.e. such that there is a world w' such that A is the complete answer to Q in w'.

So far we have not defined the notion of complete answer that we want to use, but let us first reformulate the above informal principle in a somewhat more formal way - the notion of 'complete answer' is, at this point, a parameter whose exact value we have not fixed; what counts is that the complete answer to a question in a world *w* is a proposition, i.e. has type <s,t>.

If Q is a question, let us note $Ans_Q(w)$ the complete answer to Q in w. Let P be a predicate taking both declarative and interrogative complements (i.e. P is a two-place predicate, that relates an individual to a proposition or a question). Let us call P_{decl} the variant of P that takes declarative complements (P_{decl} is of type <<s,t>,<e,t>>), and P_{int} the one that takes interrogative complements.[8] Our proposal is captured by the following *general* meaning postulate:[9]

[8] We are not committed to the view that each responsive predicate really comes in two variants in the lexicon. Rather, one has to be derived from the other by some type-shifting rule. It is also possible to define a type-shifter that would apply to the interrogative clause itself, licensing it as a complement of verbs or predicates of attitude. See Egré (2008). As is well known (Groenendijk and Stokhof 1982), responsive predicates can take as a complement a coordinate clause made up of a declarative clause and an interrogative clause (*John knows that Peter attended the concert and whether he liked it*), which suggests that the correct account should not rely on the assumption that the ambiguity is located in the responsive predicate itself. However, the choice between these various options is immaterial for everything we say in this paper, so we choose (for simplicity) to present our proposal in terms of a lexical rule (a meaning postulate) that defines the meaning of the interrogative-taking variant of a responsive predicate in terms of the meaning of its declarative-taking variant.

[9] We adopt an *intensional* semantics framework in which all expressions are evaluated with respect to a *world,* but nothing in this paper hinges on this choice. Throughout the paper, we adopt Heim & Kratzer's (1998) and von Fintel & Heim's (2011) notational conventions.

(57) $[\![P_{int}]\!]^w = \lambda Q.\lambda x_{<e>}.\ \exists w'([\![P_{decl}]\!]^w(Ans_Q(w')))(x) = 1$

According to this semantics, for John to know who came, there must be a potential complete answer ϕ to 'who came?' such that John knows ϕ. Since *know* is factive, it follows that ϕ must be true. On the other hand, for Jack and Sue to agree on who came, they must agree that ϕ, for some ϕ that is a potential complete answer to "Who came?". It is clear that ϕ does not have to be true, since *agree* is neither factive nor veridical. Generally speaking, (57) predicts that a responsive predicate is either veridical with respect to both its declarative and interrogative complements, or with respect to neither of them, which is what we argued for above. We should note here that our proposal strikingly contrasts with both Karttunen's and Groenendijk and Stokhof's proposals. To illustrate, let us restrict our attention to G&S, who proposed what amounts to the following lexical rule:

(58) $[\![P_{G\&S\text{-}int}]\!]^w = \lambda Q.\lambda x_{<e>}.\ [\![P_{decl}]\!]^w(Ans_Q(w)))(x) = 1$

This amounts to saying that, given a responsive predicate P, an individual x is in the relation $P_{G\&S\text{-}int}$ to the question Q if and only if x is in the relation P_{decl} to the actual complete answer to Q. On such an account, it is predicted that for any predicate *P*, if x is an individual in the relation $P_{G\&S\text{-}int}$ to the question Q, then x is in the relation P_{decl} to a true proposition (because the actual complete answer to Q is necessarily true). This lexical rule thus does not predict the generalization in (35), since it leads us to expect that every responsive predicate is veridical-responsive.[10] The behavior of verbs such as *tell* seemed to support this account, since it was assumed that *tell*, even though it is neither factive nor veridical when it embeds a declarative complement, is nevertheless veridical-responsive.[11] However, as we have seen, and as G&S (1993) noted themselves, verbal constructions such as *agree on* are not veridical-responsive. For this reason, G&S explicitly excluded such verbs from the domain covered by their account, which was restricted to veridical-responsive predicates (extension predicates in their terminology). It follows that in order to predict the behavior of all responsive predicates, G&S's account has to be supplemented with lexical stipulations that determine, on a case by case basis, which verbs can be subjected to the lexical rule in (58). We argued, however, that verbs such as *tell* are only apparent exceptions to the generalization (35), and this paves the way to a unified account of the semantics of responsive predicates.

Let us now turn to the exact definition of what counts as a *potential complete answer*. The proposal in (57) imposes certain constraints on what will count as a complete answer, if we are to get plausible truth-conditions for sentences in which an interrogative clauses is embedded. In particular, Groenendijk & Stokhof (1982)'s definition of complete answer is to be preferred to the weaker notion that is used in Karttunen (1977). Later in this paper we will be able to qualify this claim (see section V).

III. 1. Two notions of complete answer: weak and strong exhaustivity

In the recent literature on the semantics of interrogatives, two distinct notions of what counts as a *complete answer* to a question are usually considered. One corresponds to the so-called 'weakly exhaustive reading' of embedded questions, and is in fact the same notion as the one proposed in Karttunen (1977) (hereafter, K). The other one, which is related to the 'strongly

[10] That view is argued for in Egré (2008), where the non-veridicality of responsive predicates also taking declarative complements is attributed to the presence of overt or covert prepositions.
[11] This view is endorsed by Higginbotham (1996) in particular, who presents it as a reason not to adopt an existential semantics for questions (such as the one we endorse here).

exhaustive reading' of embedded-questions, is identical (or nearly identical) to the concept of complete answer that is assumed in theories based on 'Partition Semantics', whose most famous implementation is found in Groenendijk & Stokhof (1982, 1984, 1997, hereafter G&S). We start by showing, in a somewhat informal way, that we need to make use of G&S's notion of complete answer. In the next subsections, we provide an explicit implementation.

Consider the following question:

(59) Which students left?

According to K's definition of complete answer, the complete answer to (60) in a given world is the conjunction of all the true propositions of the form *x left*, where *x* denotes an individual who is a student in that world. According to G&S, the complete answer to the very same question in a world *w* is the proposition that consists of all the worlds *w'* satisfying the identity *student(w') ∩ left(w') = student(w) ∩ left(w)* (where, for any predicate *P* and world *v*, *P(v)* denotes the extension of *P* in *w*),[12] i.e. those worlds in which the individuals who are students in the actual world and who left are the same as in the actual world.

Suppose the relevant domain consists of three students, Mary, Susan and Ernest. Take a world *w* in which Mary and Susan left and Ernest didn't. Then the complete answer in Karttunen's sense in w is *Mary and Susan left*. But the complete answer in G&S's sense is *Mary and Susan left and Ernest didn't*. In other words, a complete answer in K's sense simply states that the students who actually came came, while in G&S's sense, the complete answer adds to this that no other student came.

Let us see what these two different notions predict when combined with (56) (or (57), which is a formal rendering of (56)):

For Karttunen, "Mary left" is the complete answer in all worlds in which Mary and nobody else left. Therefore, in the above scenario, in which Mary and Susan left, "Mary left" is, under Karttunen's definition, a *potential* complete answer, which is furthermore true in the actual world, *even though it is not the complete answer in the actual world*. It follows from (56) that if Jack knows that Mary left, then he is in the relation *know* to a potential complete answer to "who came ?". More generally, the following inference should be valid under (56):

(60) John knows that Mary left
John knows who left

[12] In fact, Groenendijk & Stokhof propose that embedded questions are ambiguous between so-called *de dicto* and *de re* readings for the restrictor of the wh-phrase. The notion of complete answer we have just defined is the one that give rise to *de dicto* readings. For the *de re* reading, one has to say, following G&S, that the complete answer to 'Which students left ?' in a world *w* is the proposition (set of worlds) which consists of all worlds *w'* such that *student(w) ∩ left(w') = student(w) ∩ left(w)*. This, however, would require that we adopt an *extensional* semantic framework in which world variables belong to the object language and occur as arguments of predicates in the logical forms being interpreted. In this paper, we are not concerned at all with the *de re/de dicto* ambiguity for restrictors of wh-phrases, so we simply ignore this ambiguity. In practice, we will only consider cases where the extension of the restrictor is common knowledge, i.e. can be thought of as constant across worlds, which makes the *de re/de dicto* distinction irrelevant.

This is of course a bad result. In fact, using Karttunen's definition of 'complete answer' in (56) would make *Jack knows who left* equivalent to *There is someone such that Jack knows that (s)he left.*[13]

For G&S, on the other hand, *Mary left* is not a potential complete answer if the domain of individuals contains at least Mary and someone else: in a world in which Mary left and nobody else did, the complete answer in G&S's sense is the proposition that states that Mary left and nobody else did. G&S's semantics for questions, when combined with (57), gives us reasonable truth conditions for *Jack knows who left*. Let us see why. What (57) tells us is that *Jack knows who left* is true if and only if there is a potential complete answer φ to *who left* such that Jack is in the relation denoted by *know* to φ. Let us assume that there exists such a proposition φ. Because *know* is factive, φ has to be true. Furthermore, there exists only one proposition that is both true and is a potential complete answer to the question, namely the actual complete answer. This is so because in G&S's semantics, complete answers are mutually exclusive, and so only one can be true in a given world. Therefore φ has to be the actual complete answer. Hence *Jack know who left* is true if and only if *Jack knows φ*, where φ is the actual complete answer to *who left ?*. As a result, this sentence is predicted to be true if and only if Jack knows that X left and nobody else did, where X is the set of all the people who left. In fact, this is exactly what G&S themselves predict.

At this point, it seems that we are forced to use G&S's notion of complete answer, since using the weaker notion gives rise to patently too weak truth conditions.

It has been argued, though, that the reading that we derive is not the only one, and that there also exists a slightly weaker reading, called the *weakly exhaustive reading*. In the case of *know*, this reading can be paraphrased as follows:

(61) *Jack knows who came* is true in w iff *Jack knows that X came*, with X being the plurality consisting of all the people who came

Consider again a situation in which Mary and Sue came and Peter and Jack didn't come, there is no other individual in the domain, and Jack knows what the domain is. Then *Jack knows who came* is true according to (61) if and only if Jack knows that Mary and Sue came. It follows that if Jack only knows that Mary came, then he does not know who came; but on the other hand, if Jack knows that Mary and Sue came but does not know that Peter didn't, then *Jack knows who came* is predicted to be true given (61). Yet Jack would not know what the complete answer is in G&S's sense. As we will see (see section V), the very existence of such a reading is under debate. So far our point is simply that our proposal cannot capture this reading, whether or not it really exists – because using the weak notion of complete answer

[13] Chemla (p.c.) suggests that this is maybe a good result in the end, as it would amount, according to him, to the *mention some* reading of wh-questions. However, we should note that things may be even worse than what we say in the main text. In fact, a formally explicit implementation of the notion of complete answer in Karttunen's framework leads to the conclusion that, in a world where no student left, the complete answer to *Who left?* is simply the tautology. This is so because in such a world there is no true elementary answer to the question, i.e. its denotation given Karttunen's semantics (i.e. the set of true elementary answers) is the empty set. Then by defining the complete answer as the conjunction of all the true elementary answers, we find that the complete answer in such a case is the tautology (because the grand conjunction of a set of propositions is the tautology when the set is empty). Karttunen introduced a special proviso for this case, in order to ensure that *Mary knows who left* be not trivially true when nobody in fact left, but rather mean that Mary knows that no one left. If our goal were to predict mention-some readings, one possible approach would be to assume that such questions introduce a presupposition that at least one elementary answer is true. We return to this issue in section V.

does not give rise to this reading, but to a much weaker one. We will return to this issue in section V.

As we shall discuss later, this may be problematic in light of recent observations that show, conclusively in our view, that *weakly exhaustive readings* exist as well, at least for some responsive predicates. In particular, *surprise*-type verbs have been argued to only support weakly exhaustive readings (Guerzoni and Sharvit 2007). We offer a solution to this problem in section V.

III. 2. Ans_2 and Ans_1

In the next sections, we will actually need to use *both* notions of complete answers. Let us make clear some of our notational choices, which come from Heim (1994):

First, we take the denotation of a given question Q, in a world w, to be a *set of propositions*, namely, those *elementary answers that are true in w*:

(62) $[\![\text{Who came?}]\!]^w = \lambda\phi_{<s,t>}.\ \phi(w) = 1\ \&\ \exists x.\ \phi = \lambda w'.x \in [\![\text{came}]\!]^w$
(informally: the set of the true propositions of the form 'x came')

Following Heim (1994), we define a first notion of complete answer to Q in w, noted $Ans_1(Q)(w)$, and corresponding to Karttunen's definition, i.e. the generalized conjunction of all the members of Q(w), which is equivalent to the following:

(63) $Ans_1(Q)(w) = \lambda w'.\forall\phi\ (\phi \in Q(w) \rightarrow \phi\ (w') = 1)$
<in words: the set of worlds in which all the members of Q(w) are true>

As we mentioned in footnote 13, in a world *w* where nobody came, *Ans₁*(*Who came?*)(*w*) is simply the tautology.[14]

Following Heim (1994), we define a second notion of complete answer, which captures G&S's notion of complete answer:

(64) $Ans_2(Q)(w) = \lambda w'.(Q(w') = Q(w))$

In other words, $Ans_2(Q)(w)$ is the proposition consisting of all the worlds in which the true elementary answers to Q are the same as in the actual world. For instance, if Q is "who came?", and if w is such that Peter, Mary and nobody else came in w, $Ans_2(Q)(w)$ denotes the proposition consisting of all the worlds in which Peter came, Mary came, and nobody else came.

Due to the above considerations, our meaning postulate deriving the meaning of P_{int} in terms of that of P_{decl} has to be based on the strong notion of complete answer, i.e. Ans_2. Thus (57) becomes:

(65) $[\![P_{int}]\!]^w = \lambda Q_{<<s,t>,t>}.\lambda x_e.\ \exists w'([\![P_{decl}]\!]^w(Ans_2(Q)(w'))(x) = 1)$

[14] This is so for the following reason. Let Q be 'Who came?'. If nobody came in *w*, we have $Q(w) = \varnothing$, and so for every world *w'*, the clause $\forall\phi\ (\phi \in Q(w) \rightarrow \phi\ (w') = 1)$ holds, with the resulst that $\lambda w'.\forall\phi\ (\phi \in Q(w)) \rightarrow \phi\ (w') = 1)$ is the constant function that maps every world to 1, i.e. is the tautological proposition.

IV. Incorporating presuppositions

Suppose now that P_{decl} is a presupposition trigger, i.e. a sentence of the form 'x P S', where x and S are not themselves presuppositional, that happens to have non-trivial presuppositions. We take this to mean that 'x P S' lacks a defined truth-value in some worlds, those in which the presuppositions of 'x P S' are not true. For instance "Jack knows that it is raining" has no truth value in a world in which it is not raining. It is a natural question whether the presuppositional behavior of P_{decl} is in some way inherited by P_{int}.

Let us have a closer look at (65). Clearly, for the condition "$[\![P_{decl}]\!]^{w}(Ans_2(Q)(w')))(x) = 1$" to hold, "$[\![P_{decl}]\!]^{w}(Ans_2(Q)(w'))(x)$" has to be defined in the first place. It follows that (65) can be reformulated as follows:

(66) $[\![P_{int}]\!]^{w} = \lambda Q_{<<s,t>,t>>}.\lambda x_e.\ \exists w'[[\![P_{decl}]\!]^{w}(Ans_2(Q)(w')))(x)$ is defined & $[\![P_{decl}]\!]^{w}(Ans_2(Q)(w')))(x) = 1$]

In more informal terms, this is equivalent to:

(67) 'x P_{int} Q' is true in w iff there is a potential complete answer S to Q such that a) the presupposition of 'x P_{decl} S' is true, and b) 'x P_{decl} S' is true.

Assume for simplicity the following lexical entry for $know_{decl}$, according to which 'x knows S' *presupposes* that S is true and *asserts* that x believes S.[15]

(68) $[\![know_{decl}]\!]^{w} = \lambda\phi.\lambda x$: $\phi(w) = 1$. x believes ϕ in w.

Then we have:

(69) $[\![know_{int}]\!]^{w} = \lambda Q_{<<s,t>,t>}.\lambda x_e.\ \exists w'$ ($[\![know_{decl}]\!]^{w}$ $(Ans_2(Q)(w')))(x)$ is defined & x believes $(Ans_2(Q)(w'))$ in w.

i.e.

(70) $[\![know_{int}]\!]^{w} = \lambda Q_{<<s,t>,t>}.\lambda x_e.\ \exists w'$ $((Ans_2(Q)(w')(w) = 1$ & x believes $(Ans_2(Q)(w'))$ in w)

Now the condition $(Ans_2(Q)(w')(w) = 1$ simply means that $Ans_2(Q)(w')$ is the actual complete answer to Q in w. So we end up with:

(71) $[\![know_{int}]\!]^{w} = \lambda Q_{<<s,t>,t>}.\lambda x_e$. there is a proposition S that is the actual complete answer to Q in w and x believes S in w

i.e.

(72) $[\![know_{int}]\!]^{w} = \lambda Q_{<<s,t>,t>}.\lambda x_e$. x believes the complete answer to Q in w

[15] The assertive meaning of *know* is in fact stronger, see Gettier (1963)'s classic arguments, but this point is not relevant for what follows.

This is in fact exactly the same as what is generally assumed in the literature. We see that the fact that $know_{int}$ is veridical-responsive follows straightforwardly from its factivity and our meaning postulate in (65)/(66).

IV.1. A presuppositional variant

So far (65) and (66) predict that even when P_{decl} is a presupposition trigger, P_{int} will not be. Rather, what follows from (65) and (66) is that for some potential complete answer S to Q, 'x P_{int} Q' *entails* (but does not presuppose) the truth of the presuppositions of 'x P S' – this is so because for 'x P S' to be true, it has to be able to have a truth value in the first place. What if we decided to turn this entailment into a presupposition? We would end up with the following (We use Heim & Kratzer's 1998 notation for representing presuppositions in our lexical entries.)[16]

(73) $[\![P_{int}]\!]^w = \lambda Q_{<<s,t>,t>}.\lambda x_e$: $\exists w'([\![P_{decl}]\!]^w(Ans_2(Q)(w'))(x)$ is defined).
$\exists w'[[\![P_{decl}]\!]^w(Ans_2(Q)(w')))(x)$ is defined & $[\![P_{decl}]\!]^w(Ans_2(Q)(w'))(x) = 1]$

In more informal terms, what (73) says is the following (from now on, when we use the phrase 'potential complete answer to Q', we mean 'a proposition S such that in some world w S is the complete answer to Q *in the strong sense*', i.e. such that S = $Ans_2(Q)(w)$):

(74) a. 'x P_{int} Q' *presupposes* that for some potential complete answer S to Q, the presupposition of 'x P_{decl} S' is true.
b. 'x P_{int} Q' *asserts* that for one such potential complete answer S, 'x P_{decl} S' is defined and true.

It turns out that even though such a modified meaning postulate normally allows the presupposition triggered by P_{decl} to be in a sense inherited by P_{int}, the verb *know* is in fact still predicted to trigger no presupposition when it embeds a question. Let us show why. Applying (74)a to *know* predicts the following presupposition for 'x knows Q':

(75) 'x knows Q' presupposes that for some potential complete answer S to Q, the presupposition of 'x knows S' is true

This is equivalent to:

(76) 'x knows Q' presupposes that for some potential complete answer S to Q, S is true

But note that the presupposition predicted by (76) is a tautology: it simply states that Q has a true complete answer, which is necessarily the case given the kind of semantics for questions we are assuming.[17]

Yet (73)/(74) and (66)/(67) are not equivalent in the general case. We shall now argue that the behavior of some other verbs that trigger more complex presuppositions, such as *agree that/agree on*, provides evidence for (73)/(74).

---

[16] Namely, a lexical entry of the form $[\![X]\!]^w = \lambda A. \ldots . \lambda Z$: $\phi(A,\ldots\ldots,Z).\psi(A,\ldots\ldots,Z)$, where *X* has a type that 'ends in t', means that *X* denotes a function which, when fed with arguments A,...,Z of appropriate types, is defined only if $\phi(A,\ldots\ldots,Z)$ is true, and, when defined, returns the value 1 if and only if $\psi(A,\ldots\ldots,Z)$ holds. In other words, the presuppositions triggered by an expression are encoded in the part for the right-hand side formula standing between the column and the period.

[17] In fact, it predicts that 'x knows Q' has the same presuppositions as Q.

IV. 2. ***Agree that/Agree on***

In this section, we will be concerned with the presuppositions and the truth-conditional content of the following types of examples:

(77) Jack agrees with Sue that it is raining
(78) Jack agrees with Sue on whether it is raining
(79) Jack agrees with Sue on which students came

(80) Jack and Sue agree that it is raining
(81) Jack and Sue agree on whether it is raining
(82) Jack and Sue agree on which students came

We will first focus on examples (77)-(79). First, we have to decide what the presuppositions of (77) are. To this end, we should see what happens when (77) is embedded under negation:

(83) Jack does not agree with Sue that it is raining

Clearly, both (77) and (83) license the inference that Sue believes that it is raining. From this we may reasonably conclude that (77) presupposes that Sue believes that it is raining, and that it furthermore asserts that Jack believes that it is raining. But this is actually insufficient. Indeed, (83) would then presuppose that Sue believes that it is raining and assert that Jack does not have this belief. This happens to be true in a situation where Sue believes that it is raining and Jack has no specific belief. But in fact, (83) seems to entail that Jack actually *disagrees* with Sue, i.e. believes that it is not raining. We can capture this fact by adding a presupposition according to which *Peter is opinionated with respect to the question whether it is raining*, i.e. Peter either believes that it is raining or believes that it is not raining. In this case, (83) will presuppose a) that Sue believes that it is raining and b) that Peter has an opinion as to whether it is raining or not, and would assert c) that Peter does not have the belief that it is raining. Together with b), c) entails that Peter believes that it is not raining, i.e. disagrees with Sue.[18]

So we end up with the following:

(84) *X agree(s) with Y that S* presupposes that Y believes that S and that X either believes S or not-S, and it asserts that X believes S.

In more formal terms:

(85) $[\![\text{agree}_{\text{decl}}]\!]^{w} = \lambda\phi_{<s,t>}.\lambda y.\lambda x{:}(Dox_y(w) \subseteq \phi)\ \&\ (Dox_x(w) \subseteq \phi \text{ or } Dox_y(w) \subseteq \neg\phi).\ Dox_x(w) \subseteq \phi$

[18] The fact that (83) entails not only that Jack does not have the belief that it is raining, but also that he believes that it is not raining, is reminiscent of the behavior of *neg-raising verbs*, such as *believe* (i.e. *Jack doesn't believe that p* is usually understood as *Jack believes that not-p*). In fact it has been proposed (Gajewski 2005) that the neg-raising behavior of *believe* should be captured in exactly the same way as what we have suggested for *agree that*: namely x *believes that P* would presuppose that x is opinionated with respect to P, and assert that x believes that P.

with $Dox_x(w)$ being defined as the (characteristic function of the) set of worlds compatible with x's belief in w (i.e. the proposition expressed by the conjunction of all of x's beliefs in w).

Let us now see what happens when we combine (85) with (73)/(74). First we consider (78), repeated below as (86), replacing *it is raining* by *p* (for convenience):

(86) Jack agrees with Sue on whether p

Regarding the semantics of whether-questions, we make the standard assumption that the denotation of *whether p* in a world w is the singleton containing the proposition expressed by p (for short: p) if p is true in w, and is the singleton containing ¬p if p is false in w. In the case of whether-questions, there is simply no difference between Ans_1 and Ans_2: both Ans_1*(whether p, w)* and Ans_2*(whether p, w)* are in fact the proposition that is the only member of the denotation of *whether p* in w. The set of potential complete answers to *whether p* is therefore simply {p, ¬p}.

(73)/(74) predicts the following for (86):

(87) *Jack agrees with Sue on whether p* presupposes that for some member S of {p, ¬p}, Sue believes S and Jack is opinionated with respect to S, and asserts that for some member S meeting this condition, Jack believes S.

This is equivalent to:

(88) *Jack agrees with Sue on whether p* presupposes that either Sue believes p and Jack is opinionated with respect to p, or Sue believes ¬p and Jack is opinionated with respect to ¬p, and asserts that if Sue believes p, then Jack believes p, and if Sue believes ¬p, then Jack believes ¬p.

Because being opiniated with respect to p is the same as being opinionated with respect to ¬p, this is in turn equivalent to:

(89) *Jack agrees with Sue on whether p* presupposes that Sue and Jack are both opinionated with respect to p, and asserts that they either both believe p or both believe ¬p.

When we turn to the negation of (86), the presuppositions remain the same, and the assertion gets negated. So we end with the following meaning for *Jack does not agree with Sue on whether p*:

(90) *Jack does not agree with Sue on whether p* presupposes that Sue and Jack are both opinionated with respect to p, and that one of them believes p and the other one believes ¬p

So far, these predictions seem to be right. Notice that the meanings given in (89) and (90) are *symmetrical* with respect to Jack and Sue. Thus, while *Jack does not agree with Sue that p* and *Sue does not agree with Jack that p* are not synonymous (the first one entails that Jack believes that not-p, and the second one presupposes that Jack believes that p), *Jack does not* *agree with Sue on whether p* and *Sue does not agree with Jack on whether p* happen to be synonymous. This seems to us to be reasonably close to what the actual facts are.[19]

Let us now turn to a more complex example, in which the embedded question is a wh-question instead of a whether-question:

(91) Jack agrees with Sue on which students came to the party

(73)/(74) results in the following:

(92) **Presupposition**: (91) presupposes that for some potential complete answer S to 'Which students came to the party', the presupposition of *Jack agrees with Sue that S* is met.
Assertion: (91) asserts that for some potential complete answer S to 'Which students came to the party', the presupposition of *Jack agrees with Sue that S* is met and it is true that Jack agrees with Sue that S.

Given our assumptions regarding the semantics of *agree with*, this is equivalent to the following:

(93) **Presupposition of (91)**: For some potential complete answer S1 to 'Which students came to the party', Sue believes S1 and Jack is opinionated with respect to S1.
Assertion of (91): For some potential complete answer S2 to 'which students came to the party', such that Sue believes S2 and Jack is opinionated with respect to S2, Jack believes S2.

It turns out that necessarily S1 and S2 are identical: indeed potential complete answers are mutually exclusive; thus since the presupposition entails that Sue believes a certain potential complete answer S1, it follows that if (91) is true, S1 is the only potential complete answer such that Sue believes it (assuming of course that Sue's beliefs are consistent). Hence (91) ends up asserting that for some potential complete answer S1, both Jack and Sue believe S1 is true. However, we predict an asymmetry between Jack and Sue at the level of presuppositions. Indeed, what (91) is predicted to presuppose is that there is a complete answer S that Sue believes and that Jack is opinionated with respect to this complete answer. But from this it does not follow that Jack himself believes any specific complete answer. For

---

[19] According to us, whatever interpretative differences can be perceived between the two examples are to be explained by the following considerations:

a) the two sentences present the same situation under different perspectives: for instance, *Jack agrees with Sue on whether p* may suggest that Sue said p and that Jack *then* agreed, i.e. that Sue's belief is in a sense "primary". Whether this should be captured in the semantics of *agree on,* rather than being an indirect effect of, say, the interactions between the subject/object status of an argument and information structure, is not clear to us. Note also that there is a use of *agree* that makes it similar to a communication verb, i.e. in which "Jack agreed with Sue" suggests that Jack said something whereby he made clear that he agrees with something that Sue said *before*, which gives rise to a temporal asymmetry between Jack and Sue. This use is not captured by our semantics
b) We made the assumption that *Jack agrees with Sue that p* presupposes that Jack is opinionated w.r.t.p. This was meant to capture the "neg-raising" behavior of *X does not agree with Y that p*. It might be that this presupposition is sometimes not present, in the same way as *believe* is not systematically neg-raising. If we remove this presupposition, then *Jack doesn't agree with Sue on whether p* is not predicted to have exactly the same presuppositions as *Sue doesn't agree with Jack on whether p*; the first one could be felicitous and true if Sue believes it is raining and Jack is uncertain, but not the second one.

Jack could believe that S is in fact false, thus being opinionated with respect to S, without believing any other specific potential complete answer to be true. It is hard to assess whether this is a correct prediction, because (91) anyway *entails* that there is a complete answer that Jack believes (even though this is not predicted to be *presupposed*).

But the facts become much clearer when we turn to the negative counterpart of (91), namely (94). In this case we end up with the same presupposition as before, and the whole meaning is given in (95).

(94) Jack does not agree with Sue on who came to the party.

(95) **Presupposition**: for some potential complete answer S1, Sue believes S1 and Jack is opinionated about S1.
Assertion: for some potential complete answer S2 such that Sue believes S2 and Jack is opinionated about S2, Jack does not believe that S2 is true.

Again, S1 and S2 are necessarily identical. Given that it is presupposed that Jack is opinionated with respect to S1, and that it is asserted that Jack does not believe S1, it follows that Jack believes that S1 is false. So the end result is that (94) conveys the information that for some potential complete answer S such that Sue believes S to be true, Jack believes that S is false.

Note that from this it does not follow that Jack has a definite opinion as to what the correct answer to 'which students came the party' might be; he must believe that a certain potential complete answer is false, but he might be uncertain about the truth of all the other ones. On the other hand, Sue has to have a definite opinion. So a clear asymmetry is predicted. This result seems to us to be exactly correct.

Let us now consider the 'reciprocal' use of *agree* as exemplified in (80) - (82), repeated below as (96) - (98)

(96) Jack and Sue agree that it is raining.
(97) Jack and Sue agree on whether it is raining.
(98) Jack and Sue agree on which students came.

We assume that a sentence of the form *X and Y agree that p* (instantiated by (96)) presupposes that both X and Y are opinionated as to whether it is raining, and assert that both believe that it is in fact raining. Since the presupposition and the assertion are fully symmetrical with respect to X and Y, we don't expect any asymmetry when reciprocal *agree* takes an interrogative clause. For (97), we predict that the presupposition should be that there is a complete answer *A* to *is it raining?* such that both Jack and Sue have an opinion regarding A's truth-value. This is the same as saying that both Jack and Sue have an opinion as to what the correct answer to *is it raining?* might be. And the assertion simply states Jack and Sue have the same opinion. Things get more complex for (98). In this case, we predict the following:

(99) **Presupposition:** for some potential complete answer A to *Which students came?*, both Jack and Sue are opinionated with respect to A.
Assertion: for some potential complete answer A to *Which students came?*, both Jack and Sue believe that A is true.

This seems reasonable, but in order to tease apart the presuppositional component and the assertive component, we should look at the negated version of this sentence.

(100) Jack and Sue do not agree on which students came

The predictions for (100) are as follows:

(101) **Presupposition:** there is a complete answer A to *Which students came?* such that both Jack and Sue are opinionated with respect to A.
Assertion: there is a complete answer A to *Which students came?* such that both Jack and Sue are opinionated with respect to A, and either Jack or Sue believes that A is true and the other one does not have the belief that A is true.

This amounts to saying that one of the two has a definite view as to what the complete answer to *Which students came?* might be, and that the other one does not share this view. Importantly, then, we do not predict (100) to imply that both Jack and Sue have a complete belief as to which students came to the party, but rather that only of them does. That is, (100) is expected to be true and felicitous in a situation where, say, Sue believes that all the First year students came and nobody else did, while Jack does not have a clear idea about which students came but believes that it is not true that all the First year students came and nobody else did. On the other hand, in a situation in which neither Jack nor Sue have a complete idea of which students came but happen to have a disagreement about the students for whom they do have an opinion, (100) should fail to have a defined truth-value. Whether these fine-grained predictions are correct is not clear to us. We believe it could be assessed only by gathering controlled data by experimental means.

V. Weak and Strong Exhaustivity

In this final section, we propose a modification of our basic proposal, so as to be able to capture so-called weakly exhaustive readings. We will have nothing to say about the fact that certain predicates admit both the weakly and the strongly exhaustive readings, while others have been said to be compatible with only one of them (George 2011 and Klinedinst & Rothschild claim that *know* is only compatible with the strongly exhaustive reading, while according to Guerzoni & Sharvit 2007 *know* licenses both the weakly and the strongly exhaustive reading, but *surprise* only supports the weakly exhaustive reading). We have a more modest goal here: we want to formulate a second *general* meaning postulate that turns a predicate that takes declarative complements into one that takes interrogative complements, and we want this meaning postulate to generate weakly-exhaustive readings. We assume that some additional principles that we don't investigate here sometimes eliminate one of the two readings, depending on specific properties of the predicate – see Guerzoni 2007 for ideas along these lines.

First we present some preliminary remarks about the distinction between weakly exhaustive and strongly exhaustive readings.

V. 1. Does *know* + *Q* license a weakly exhaustive reading?

Consider the following sentence:

(102) Jack knows who came.

As explained in section III.1, since Groenendijk and Stokhof (1982), there has been a debate as to whether the truth of this sentence simply implies that Jack knows of everyone who came that they came (see Karttunen's notion of complete answer), or whether it also implies that Jack knows of everyone who did not come that they did not (viz. Groenendijk and Stokhof's notion of complete answer). While everybody seems to agree that the strongly exhaustive reading exists, there is no agreement as to whether it is the *only* reading (see Groenendijk and Stokhof 1982 and more recently George 2011 for arguments against the existence of a weakly exhaustive reading, and Guerzoni & Sharvit 2007 for arguments for). One argument for the existence of the weak reading is based on the fact that the following is not deviant (Guerzoni & Sharvit 2007):

(103) Jack knows who came but he does not know who did not come

Guerzoni & Sharvit (2007) note that, given a certain fixed domain of individuals D, under the strongly exhaustive reading, 'Jack knows who came' is equivalent to 'Jack knows who did not come'. Indeed, let X be the people who came and let Y be the people who did not come. Then both *Jack knows who came* and *Jack knows who did not come*, under the strong reading, are predicted to mean that Jack knows that everyone in X came and that everyone in Y did not come. Hence (103) should be a contradiction, and the fact that it isn't is evidence for the existence of another reading.

Yet, as G&S had already pointed out, this reasoning is clearly dependent on there being a constant domain of inviduals (see also George 2011). For suppose that the domain of quantification is not the same in every world. Suppose that in the actual world, X are the people who came and Y are the people who did not come. Then *Jack knows who came* is predicted to mean that Jack knows that X came and that nobody else did. But from this it does not follow that Jack knows that Y did not come and that every person not in Y came, since Jack might not know which individuals there are in the actual world, hence fail to know the extension of Y. He may know that nobody apart from John came, and yet have no idea whatsoever who did not come (maybe he does not even know whether there are people who are not John).

In fact, G&S make the prediction that (103) is a contradiction only if the domain of quantification is kept constant across all the epistemic alternatives of Jack. So the question is whether when we consider a scenario where this has to the case, (103) starts sounding like a contradiction. That is, does (104) sound contradictory?

(104) Mary has just shown Jack the list of the people who were invited to the party, and she wants to ask him which of these people actually came.
Jack knows who came to the party, but he does not know who did not.

While it seems to us that (104) sounds contradictory, we should not be too quick to jump to the conclusion that the weakly exhaustive reading does not exist for *know*, for it could be that the kind of context used here biases the interpretation towards the strongly exhaustive reading, even if the weakly exhaustive reading were in principle available.[20]

[20] See George (2011) for an extensive discussion. George concludes that the weakly exhaustive reading does not exist for know. In order to reach a firm conclusion, one should probably carry out a systematic survey on naïve informants (see Rothschild & Klinedinst 2011 report the results of such a survey for the verb *predict*, but not *know* – see sections V.4 and V.5).

However, it is important to point out the following fact, which has often not been recognized: even if some kind of weakly exhaustive reading exists for *know*, its standard characterization cannot be right, for reasons discussed in Spector (2005, 2006), and George (2011). To see what the point is, consider the following situation. Suppose that, among a number of people invited to a party, the guests who attented the party are exactly John, Peter and Sue. Now suppose that Mary knows that John, Peter and Sue attended, but also believes, wrongly, that Al attended. In such a case Mary *is* in the relation denoted by *know* to the complete answer in Karttunen's sense. And yet it seems to us and all our informants that she cannot be said, on any conceivable reading, to *know which of the guests attended the party*.[21]The fact that Mary has a *false* belief that someone who in fact did not attend the party did attend the party makes the sentence *Mary knows which of the guests attended the party* clearly false. Yet the weakly exhaustive reading as usually defined would make the sentence *Mary knows which of the guests attended the party* true in this scenario. We can thus conclude that, in any case, we don't want our theory to generate weakly exhaustive readings in the standard sense.

Now, as pointed out in Spector (2005, 2006), there is a way of characterizing the weakly exhaustive reading that does not run into this problem (see also George 2011 and Klinedinst & Rothschild 2011). Namely, one could say that *X knows Q* is true if a) X knows the truth of what is in fact the complete answer in Karttunen's sense, and b) there is no *stronger* potential complete answer in Karttunen's sense that X *believes*. According to such a semantics, *Mary knows which of the guests attended the party* cannot be true in the above scenario, because Mary believes the proposition 'John, Peter, Sue and Al attended', which asymmetrically entails the actual complete answer in Karttunen's sense. As pointed out in George (2011), however, such a meaning cannot be derived in the classical approaches to embedded questions. With the exception of George (2011), all current proposals, just like ours, assume that the truth-value of *X knows Q* only depends on the identity of the propositions to which X is related via the attitude relation *know*. But according to the proposed meaning for *X knows Q* just described, the truth-value of *X knows Q* would depend not only on what X knows, but also on what X *believes* (cf. also Klinedinst & Rothschild 2011) without knowing (such as false beliefs). To conclude, our approach so far is not is able to derive the weakly exhaustive that other approaches derive, but this does not seem to be very problematic in the case of *know*, since the weakly exhaustive reading as usually characterized does not seem to correspond to an attested reading. There may exist, nevertheless, a slightly different 'weakly exhaustive reading' for *know* + *question* (the one we have just described), but this reading is beyond the reach of current approaches, unless they are substantially modified. We will briefly return to this issue in section V.4. In the next sections we will show that we need to be able to derive weakly exhaustive readings for at least some embedding verbs, and we will propose a modification of our proposal that does capture such readings.

V.2. *surprise* and weakly exhaustive readings.

While it is difficult to reach a firm conclusion in the case of *know,* a case can be made that the weakly exhaustive reading does exist when we turn to some other predicates. *Surprise* is a case in point. As argued by Guerzoni & Sharvit, embedded questions under *surprise* and other so-called emotive predicates (*amaze*) favor a weakly exhaustive reading. Consider for instance the following sentence.

(105) It surprised Mary which of our guests showed up.

---

[21] Lahiri (2002) briefly mentions a similar example, which he attributes to J. Higginbotham. Namely, the sentence *John knows which numbers between 10 and 20* are prime cannot be true (on any conceivable reading) if John happens to believe that *all* numbers between 10 and 20 are prime.

Consider the following scenario. We had a party in which we invited only married couples. Mary expected three couples, let's call them A, B, and C, to show up, and didn't specifically expect any other guest to show up. What she definitely did not expect is that a guest would show up without his or her spouse. In fact, it turned out that only one person in each couple A, B and C showed up. She did expect these people to come, but she did not expect that their spouses would not. Guerzoni & Sharvit's point is that in such a scenario, (105) does not seem to be true. While it surprised Mary which of the guests did *not* show up, it did not surprise her that the guests who actually showed up showed up. This provides an argument that *surprise*, when it embeds a constituent question, only licenses the weakly exhaustive reading.[22] That is, (105) is true if and only if there is a complete answer in K's sense, i.e. a proposition of the form *X showed up*, where *X* is a plurality of guests, such that the truth of this proposition surprised Mary.

As we have seen, our proposal so far is not able to generate this reading. However, we will now see how we can slightly amend it in order to correctly predict the truth conditions of sentences such as (105).

V.3. Generating weakly exhaustive readings within our framework

So far, the reason why we were not able to generate the weakly exhaustive reading for embedded interrogatives is the following: if we use the strong notion of complete answer (i.e. Ans_2), we get the strongly exhaustive reading. But if we use the weak notion, as shown in section III.1, we don't get the weakly exhaustive reading, but a much weaker reading. In footnote 13, we noted that this weaker reading might in fact be the *mention-some* reading of questions, but we also noted problems with this view.

However, in the subsequent sections, we made our lexical rule for the derivation of question-taking predicates more complex, by taking into account certain facts about presupposition projection. It turns out that once we allow ourselves to formulate lexical rules that have both a presuppositional and a non-presuppositional component, we are in fact in a position to propose a second lexical rule that does capture the weakly exhaustive reading. We will start from an examination of the behavior of emotive-factive predicates for which there is widespread argreement that they license weakly exhaustive readings. We will then discuss how our modified account can deal with weakly exhaustive readings with *predict* and *know*, assuming those can exist.

<u>V.3.1 *Surprise+interrogative*: strong exhaustivity in the presupposition, weak exhaustivity in the assertion.</u>

The main idea behind this second lexical rule is the following: in its presuppositional part, the rule makes reference to the complete answer in the strong sense, while in the assertion part, the rules makes reference to both the strong notion of complete answer and the weak notion. Let us illustrate how this could work at an informal level, in the specific case of *surprise*. First, note that declarative-taking *surprise* triggers both the presupposition that its complement is true and that the attitude holder knows that the complement clause is true.

(106) It surprises Mary that p
→ Presupposition: p and Mary believes p (= Mary knows p)

[22] Rothschild & Klinedinst (2011) argue that *surprise* can in fact give rise to strongly exhaustive readings. The important point for us is only the fact that a weakly exhaustive reading is clearly available (and seems in fact to be preferred).

Let us call the complete answer to a question in Karttunen's sense the *K-complete answer* (or weakly exhaustive answer) and the complete answer to a question in G&S's sense the *G&S-complete* answer (of strongly exhaustive answer). As explained above, the *K-complete answer* to a question *Q* in a world *w* is $Ans_1(Q)(w)$, and the *G&S-complete answer* to *Q* in *w* is $Ans_2(Q)(w)$. What our proposal will derive for *It surprises Mary Q*, where *Q* is a constituent question, is the following:

(107) **Presupposition:** there is a world w such that the G&S-complete answer to Q in w, call it A, is such that the *presupposition* of *It surprises Mary that A* is true.
Assertion: there is a world w such that a) the G&S-complete answer to Q in w, call it A, is such that the *presupposition* of *It surprises Mary that A* is true and b) the K-complete answer to Q in w, call it A', is such that it surprises Mary that A'.

Let us see what this amounts to in the case of (105). The predicted presupposition is that there is a potential G&S-complete answer A to 'Which guests showed up?' such that the presuppositions of *it surprises Mary that A* are true, i.e. such that A is true and Mary knows that A is true. Now, because the only potential G&S-complete answer that is true in the actual world is the actual G&S-complete answer, this amounts to saying that the actual G&S-complete answer is true (a tautology) and that Mary knows that it's true. Because G&S-complete answers are not mutually compatible, if Mary knows that one is true, then if she is coherent she also knows that it is the actual complete answer. Hence the predicted presupposition amounts to: *Mary knows what the G&S-complete answer is*, i.e. *Mary knows which of the guests showed up*, in the strongly exhaustive sense.

Let us now turn to the assertive part. The assertive part says that there is a world w such that the G&S-complete answer in w is true, Mary knows that it is true, and such that the truth of the K-complete answer surprises Mary. Now, again, the world that possibly can satisfy all these conditions is the actual world. So the assertion can be reformulated as follows: Mary knows the truth of the actual G&S-complete answer, and she is surprised by the actual K-complete answer. This entails that Mary is surprised by the "positive" part of the G&S-complete answer, i.e. by the fact that a certain group of people came. So if Mary's only cause for surprise has to do with the guests who did **not** show up, the sentence (105) doesn't end up being true, which accounts for Guerzoni & Sharvit's observation. Note, however, that we *do* predict a very strong presupposition, according to which Mary knows which guests showed up in the strongly exhaustive sense.

Is this prediction correct? The facts seem to us to be rather complex, but we think that once we properly control for the quantificational domain, the data go the way we predict. Let us see this. In order to make sure that the domain of quantification for the wh-phrase is explicitly fixed, let us imagine the following sentence in the specified context:

(108) <u>Context:</u> Mary has 10 students, and they all took a certain exam. She definitely did not expect students A, B and C to pass, she had no specific expectations for others. In fact, students A, B, C passed and no other student did. A, B and C sent her an e-mail to tell her that they passed. She was surprised. Regarding the seven other students, she has no information, i.e. does not know yet whether they passed or not (even though in fact they didn't pass). Now, John and Sue know all this, i.e. they

know both which students passed and which didn't, what Mary knows and doesn't know, and are aware that she was surprised that A, B and C passed. In fact, they overheard her saying "I can't believe that A, B and C passed! As to the other 7 students, I don't know yet whether they passed". John and Sue are looking at a list of the ten students, and John then tells Sue the following:

Sentence: "It surprised Mary which of her 10 students passed".

According to the 5 native English speaker we consulted, the sentence sounds awkward, precisely because John and Sue *know* that Mary doesn't know exactly which of her 10 students passed.[23]That is, even though Mary is surprised by what is, *in fact*, the K-complete answer to "Which of Mary's ten students passed", the fact that she does not know that this proposition *is* the complete answer (because she cannot exclude so far that other students passed) creates a presuppositional failure for the sentence.

V.3.2. Formal implementation

Let us now see how we can turn our specific proposal for surprise into a general lexical rule giving rise to a weakly exhaustive reading *at the assertive level only*. Our proposal is the following:

(109) Let P be a predicate that can take a declarative clause as a complement. Then the interrogative-weakly-exhaustive variant of P (if it exists), noted $P_{K\text{-int}}$, has the following lexical entry:

$[\![P_{K\text{-int}}]\!]^w = \lambda Q.\lambda x : \exists w' [\![P]\!]^w(Ans_2(Q)(w'))(x)$ is defined.
$\exists w' [\![P]\!]^w(Ans_2(Q)(w'))(x)$ is defined & $[\![P]\!]^w(Ans_1\text{-}Q(w'))(x) = 1]$

One can easily see that our semantics for *surprise+interrogative* is an instantiation of this rule. The trick we use in order to ensure that using the weak notion of complete answer in the assertion does not give rise to too weak results (cf. our discussion in section III.1)[24] is to make

---

[23] We contrasted the above scenario with a minimally different one in which Mary knows that A, B and C passed and that no other student passed. Our informants then had no problem with "It surprised Mary which of here 10 students passed". All our informants also reported a clear contrast between the two following discourses:

(i) # John learnt that Peter and Sue passed. Mary and Alfred failed, but John didn't learn that. It surprised him which of the four students passed.
(ii) John knows that Peter and Sue passed and that Mary and Alfred failed. It surprised him which of the four students passed

[24] One could wonder whether a very weak semantics in terms of K-complete answers, of the type we rejected in section III.1, could not be sufficient for *surprise+wh-question*. According to such a semantics, *It surprised Mary which guests showed up* would count as true as soon as there is a guest or a plurality of guests who showed up and such that the facts that this and these guests showed up surprised Mary. This is in fact a suggestion that Ben George (2011) made. However, it seems to us that such a proposal is too weak, for at least two reasons. First, imagine that Ann only knows that Mary attended a certain party, and does not know that Jane and Sue did as well. Assume further that it surprised Ann that Mary attended the party. In such a situation, a very weak semantics of the form we rejected in section III.1 predicts that the sentence *It surprised Ann which of the guests attended the party* shoul be felicitous and true. But according to our informants, for the sentence to be felicitous, Ann has to know which guests attended the party. Second, there may be situations where Mary could be surprised that a specific guest showed up and yet fail to be surprised by which guests showed up. Imagine for instance the following situation: the fact that Peter showed up is surprising to Mary, but that the overall list of

sure that the potential K-complete answer whose existence can make true the existential statement "There is a potential K-complete answer A such that x is related to A via the attitude relation P" corresponds to a G&S-complete answer A* which is such that the presuppositions of 'x V A*' are true. One can reformulate (109) as follows, so as to make this idea clearer.

(110) a. Auxiliary Definition
Given a question Q, and a potential K-complete answer A to Q, we define the G&S-associate of A as the proposition A^{*Q} that states that A is the actual K-complete answer to Q.
In other words, we have: $A^{*Q} = \lambda w'.(Ans_1(Q)(w')=A)$.
Note that A^{*Q} is necessarily a potential G&S-complete answer. The operator ^{*Q}, when applied to a potential K-complete answer, is simply an exhaustivity operator that turns it into its 'associated' G&S-complete answer.[25]
b. Lexical Rule for the weakly exhaustive reading
$[\![P_{K\text{-int}}]\!]^w = \lambda Q.\lambda x$: There is a potential K-complete answer to Q, call it A, such that $[\![P]\!]^w(A^{*Q})(x)$ is defined. There is a potential K-complete answer to Q, call it A, such that $[\![P]\!]^w(A^{*Q})(x)$ is defined and such that $[\![P]\!]^w(A)(x) = 1$.

V.4. Deriving the weakly exhaustive reading for *know* (and other factive predicates)

We now show that by applying this second lexical rule to *know*, we predict the standard weakly exhaustive reading. Using the formulation given in (110), we get the following, where *w* is the world of evaluation.

(111) Mary knows which guests attended the party
a. **Presupposition**: there is a potential K-complete answer A to 'which guests attended the party?' such that $[\![know]\!]^w(A^{*Q})(Mary)$ is defined, i.e. such that A^{*Q} is true.
b. **Assertion**: there is a potential K-complete answer A to 'which guests attended the party?" such that a) A^{*Q} is true and b) $[\![know]\!]^w(A)(Mary) = 1$

As before, the presupposition is trivial: it just says that there is a potential complete answer in the Karttunen sense such that its associated complete answer in the G&S sense is true. This is always the case, since the K-complete answer in the world of evaluation is such that its associated G&S-complete answer is true (is in fact the unique true G&S-complete answer). Let us now turn to the assertion. First, note that for any potential G&S-complete answer B to a given question Q, there is a unique potential K-complete answer A such that B = A^{*Q}, and this K-answer is the actual K-complete answer if and only if B is true.[26] There is thus a natural

guests who showed up is not so surprising, despite Peter's presence. In such a situation it is conceivable that one could say that it surprised Mary that Peter showed up, but not that it surprised her which guests showed up. Whether or not our intuition is robust, it is very possible that *surprise* is non-monotonic with respect to its complement, and we do not want our theory of question embedding to be dependent on too specific assumptions about the lexical semantics of *surprise* (such as the assumption that *it surprised Mary that p* entails *it surprised Mary that p & q*, which seems natural at first sight but might not be correct after closer scrutiny).

[25] There is of course a close relationship between '^{*Q}' and the operator Ans_2.
Namely, we have, for any Q, w : $Ans_2(Q)(w) = [Ans_1(Q)(w)]^{*Q}$.

[26] Proof: let B be a potential G&S-answer to Q. Consider a world *w* in which B is true. Let A be the actual K-complete answer to Q in w. By definition of '^{*Q}', A^{*Q} is necessarily true in *w*. Since there is only one true G&S-complete answer in *w*, and since B is true in *w*, necessarily B = A^{*Q}. Furthermore, A is the actual K-complete answer to Q if and only if A^{*Q} is true, i.e. if and only if B is true.

mapping betwen potential K-complete answers and potential G&S-complete answers. Thanks to this fact, the assertion ends up meaning the following: there is a potential G&S-complete answer that is true and such that its corresponding K-complete answer is known by Mary. Now, in a given world *w* there is only one true potential G&S-answer, and its corresponding K-complete answer is the actual K-complete answer. So the assertion ends up equivalent to 'Mary knows the truth of what is in fact the complete K-asnwer'. Since the presupposition, as we have just seen, is trivial, this is just the same as the weakly exhaustive reading as standardly defined.

So how does our proposal fare? We managed to solve the problem that we noted in section III, which prevented us from using Karttunen's notion of complete answer. The problem we had then was that we predicted too weak truth-conditions (*X knows Q* was predicted to be true as soon as X knows a correct *partial* answer to Q, possibly even the tautology). By making use of *both* the strong notion and the weak notion of complete answer in our second lexical rule, we managed to solve this problem. However, we pointed out in section V.1 that it is not clear whether the weakly exhaustive reading exists and, more importantly, that even if it exists the truth-conditions predicted by standard theories of weakly exhaustive answers are anyway too weak, for they do not include what George (2011) dubs the 'no-false-belief' constraint. That is, for *X* to *know which guests attended the party*, X should clearly not have the false belief that a certain guest who in fact didn't attend the party attended the party. We are not aware of any fully satisfying account of the weakly exhaustive reading that would solve this problem without stipulations specifically tailored to solve it. We can however suggest a solution along the lines of Klinedinst & Rothschild (2011). Adapted to our own system, their proposal would take the following form. First, we strengthen the lexical rule for the weak exhaustive reading as follows (the added condition is underlined).

(112) **Lexical Rule for the weakly exhaustive reading**
$[[P_{K\text{-}int}]]^w = \lambda Q.\lambda x$: There is a potential K-complete answer to Q, call it A, such that $[[P]]^w(A^{*Q})(x)$ is defined. There is a potential K-complete answer to Q, call it A, such that $[[P]]^w(A^{*Q})(x)$ is defined and such that $[[P]]^w(A)(x) = 1$ <u>and such that there is no potential K-complete answer to Q, call it B, such that both B asymmetrically entails A and $[[P]]^w(B)(x) = 1$.</u>[27]

At the level of the presupposition, nothing would change. But at the level of the assertion, instead of saying that *John V Q* is true if John is in the relation V to some potential K-complete answer to Q, we now say that it is true if a) John is in the relation V to some potential K-complete answer A to the question Q, and b) that there is no other potential K-complete answer B to Q such that B asymmetricaly entails A and John is the relation V to B. This move, however, is not sufficient. As Klinedinst & Rothschild (2011) themselves note, for this to work in the case of *know*, we would need to assume that the underlined clause treats *know* as synonymous with *believe*. This is so because otherwise the underlined clause would be vacuously true. Indeed, since A is necessarily the actual K-complete answer (as shown above), any potential K-complete answer B such that B asymmetrically entails A is necessarily false, with the result that John cannot be in the relation *know* to B, even if he *believes* B. But the meaning we want to derive is one that excludes that John *believes* a false potential K-complete answer. What Klinedinst & Rothschild (2011) briefly suggest is that, to

[27] Rothschild & Klinedinst derive the underlined clause indirectly, by applying an exhaustivity operator to the V+Q constituent. The underlined clause is indeed very similar to the content of an exhaustivity operator, which, when applied to a proposition A, returns A with the conjunction of A's alternatives that are not entailed by A.

the extent that the weakly exhaustive reading exists for *know* (in fact they call it the 'intermediate reading' to distinguish it from the reading that is predicted by standard theories), it shows that the lexical representation for *know* distinguishes between the presuppositional, factive component (truth of the complement), and its assertive component (whatever it is exactly). This might require a bi-dimensional theory of presuppositions, as originally proposed by Karttunen & Peters (1979). With such a tool, one could ensure that the lexical rule in (112) makes reference to the presuppositional dimension in its first part (the part that uses the strong notion of complete answers) and to the assertive dimension in its second part. We will not provide a formal implementation of this idea in this paper.[28]

V.5. The case of veridical-responsive *predict*

Klinedinst & Rothschild (2011) specificaly argue, based on the results of a systematic qustionnaire, that *predict*, in its responsive-veridical reading, licenses the 'intermediate reading', i.e. a reading where *X predicts +wh-Q* is true if and only if *X predicted the true complete answer in K's sense to wh-Q and did not predict any stronger potential complete answer*. They can derive this reading straightforwardly, because in their system the veridical-responsive reading for *predict* can be derived from a non-factive lexical entry for *predict+declarative clause*, so that the complication we have just noted for *know* does not arise.[29] However, such a move is not available for us. In order to generate a veridical-

---

[28] Two remarks are in order:

a) In the case of *agree on*, the predictions made on the basis of (109)-(110) and (112) and the data that are needed to evaluate them are quite complex, and we do not discuss them in this paper. Our impression, however, is that in order to get a plausible weakly exhaustive reading for *agree on*, we would need again to use (112) and to be able to 'extract' an assertive component of *agree* (namely *believe*, on the assumption that *X agrees with Y that p* presupposes that Y believes p and that X is opinionated w.r.t to p, and asserts that X believes p) that would be relevant to the underlined clause of (112).

b) If applied to *surprise*, the lexical rule proposed in (112), together with the view that the underlined clause uses only the 'assertive' component of *surprise*, leads to clearly wrong predictions. Namely, for *it surprised Mary which guests attended the party*, we we would get the following truth-conditions: 'it surprised Mary that the people who actually came came, but she would not have been surprised if other people had come on top of the people who actually came', which seems to us to be a clearly wrong result, possibly even a contradiction if *surprise* is monotone decreasing with respect to its complement. Klinedinst & Rothschild (2011) do not discuss this case, but their account would not run into this problem if *surprise* is treated as montone-decreasing with respect to its complement. In their system, the content of the underlined clause in (112) corresponds to the contribution of an exhaustivity operator scoping over the verb. Because the exhaustivity operator is defined in terms of entailment, it is sensitive to the monotonicity properties of its prejacent. Because *X is surprised by Y* is entailed by *X is surprised by W* when Y entails W, the presence of an exhaustivity operator in *It surprised Mary + Question* would not lead to the negation of 'It surprised Mary that W', where W is a potential complete answer to Q in K's sense that entails the actual complete answer in K's sense. In order to solve this problem, we could modify our lexical rule in (112) so as to make it sensitive to the monotonicity property of the embedding verb, by incorporating the contribution of an exhaustivity operator scoping about the verb into the lexical rule itself. However, Klinedinst & Rothschild (2011), as well as a modification of (112) along the lines we have just suggested, may run into another problem. Namely, 'It surprised Mary which guests attended the party' might end up meaning 'it surprised Mary that the guests who actually attended the party attended the party, but for every proper subgroup X of the guests who actually attended, she is not surprised that X attended the party.'

[29] In Rothschild & Klinedinst (2011), the truth-conditions of *X V Q* are not derived by existentially quantifying over potential complete answers. Following previous works, they assume that *X V Q* means '*X is the relation V to the actual complete answer to Q'*. As we discussed at length, this generates a veridical-responsive reading even for predicates which are not veridical with respect ot their declarative complements, and thus does cannot account for the full range of facts. The intermediate reading they generate for *X predicted wh-Q* is thus '*X predicted the actual complete answer in K's sense to wh-Q, and didn't predict any stronger potential complete answer*'. Because *both* occurrences of *predict* in this informal paraphrase are non-veridical, there is no need in this case to make any distinction between the presuppositional part of *predict* and its assertive part. Such a move

responsive reading for *predict*, we need to start from a factive lexical entry for *predict* (and we argued that *predict* is in fact ambiguous between a factive and a non-factive use), so that we would need, just as is the case for *know*, to be able to 'extract' from this factive lexical entry a non-factive assertive part which would be the one relevant to the underlined clause of (112).

While these remarks are fairly speculative, we hope to have shown that our proposed account has at least as much empirical coverage than other existing accounts, and at the same time meets the desideratum that we mentioned in the begining of this paper, namely provides *uniform* mechanisms to derive the meaning of 'V+Question' construction in terms of the meaning of the 'V+declarative construction'.

V.6 No weak exhaustivity for non-veridical *predict* (and other non-presuppositional responsive predicates)

When applied with *non-presuppositional* predicates, the lexical rules we proposed in (109) in (112), remain too weak, because the aspects of this rules that make reference to the presuppositional content of the relevant embedding predicate are vacuous in such a case. Take for instance non-veridical *predict*. For *Mary predicted which guests attended the party*, we get the following on the basis of (112):

(113) Mary predicted which guests attended the party.
Presupposition: there is a potential K-complete answer A to 'which guests attended the party?' such that $[[\text{predict}]]^w(A^{*Q})(\text{Mary})$ is defined.
→ This is always true (because non-factive *predict* has no presupposition)
Assertion: there is a potential K-complete answer A to 'which guests attended the party?" such that $[[\text{predict}]]^w(A)(\text{Mary}) = 1$ <u>and there is no potential K-complete answer B such that B asymmetrically entails A and $[[\text{predict}]]^w(A)(\text{Mary}) = 1$</u>.

The pressupposition is vacuous. As to the assertion, it is true as soon as Mary predicted that *some* potential K-complete answer is true. Take indeed the *strongest* potential K-complete answer that Mary predicted. Such a proposition satisfies the existential statement of the first part of the rule and satisfies the underlined clause as well. In fact, since the tautology itself *is* a potential K-complete answer (as we noted in footnote 13 and in section III.2), the resulting truth-conditions are simply that *Mary predicted which guests attended the party* comes out true as soon as she predicted the tautology, which we may assume is itself a tautology.

Given this result, we assume that using the lexical rules in (109) or (112) is ruled out pragmatically in the case of non-veridical *predict,* because it gives rise to a trivially true proposition. This point can be generalized to all other non-presuppositional predicates *P* such that *X P + tautology* is itself a tautology (presumably non-veridical *tell*, e.g.).

We therefore expect that non-veridical *predict* should *only* give rise to a strongly exhaustive reading, since for non-veridical *predict* only the lexical rule based on G&S-complete answers (the rule in (73)) is applicable. This expectation seems to us to be clearly borne out. To see this, consider the following sentence in the specified scenario:

---

is not available to us. In our own framework, using the *non-veridical* lexical entry for *predict* together with the lexical rule (112) would yield much too weak truth-conditions, for the same reasons as those discussed in section III.1.

(114) a. Scenario: Peter wondered who would attend a certain party among four students, John, Sue, Fred, and Mary. He predicted that Sue would go and made no prediction regarding the others. In fact it turned out that Sue did not attend the party
b. Sentence: Peter predicted which of the four students would attend the party, but he proved wrong

It is clear, we believe, that (114)b is false in such a scenario. We can contrast this scenario with another one where Peter made a *complete*, *exhaustive* prediction that turned out to be false.

(115) a. Scenario: Peter wondered who would attend a certain party among four students, John, Sue, Fred, and Mary. He predicted that John and Sue would go and that Fred and Mary would not. In fact the reverse turned out to be true: only Fred and Mary attended the party.
b. Sentence: Peter predicted which of the four students would attend the party, but he proved wrong

Even though the non-veridical use of *predict* is somewhat dispreferred for some speakers, there is a very clear contrast between (114) and (115), in that and (115), unlike (114), can be judged true. This shows that the non-veridical use of *predict*, which may be less salient that its veridical use, is in any case only compatible with the strongly exhaustive reading.

VI. Conclusion

In this paper, we have offered a unified theory that generates the meaning of responsive predicates on the basis of the meaning of their 'declarative' counterparts. The starting point of our theory is fairly simple: if *V* is a verb that takes both a declarative clause and an interrogative clause as a complement, the truth-conditions of *X V Q* are just '*X is in the relation V to some complete answer to Q*'. We saw, however, that this simple idea needs to be implemented in a very careful way in order to account for the full range of facts.

First, we defended this approach against a potential objection based on 'communication verbs'. We then gave an in-depth examination of the way the presuppositions induced by a verb relative to its declarative complement are 'inherited' when the verb embeds an interrogative clause. This allowed us to define two general lexical rules, one based on the strong notion of complete answer as defined in G&S, and another one that uses both this strong notion and the weaker notion of complete answer, based on Karttunen's semantics for interrogatives. This puts us in a position to derive the so-called *weakly exhaustive reading* for emotive factive verbs such as *surprise*, and to make a striking prediction, namely that the presuppositional content of a sentence of the form *It surprised X Q* is that *X knows the complete answer to Q* in the strongly exhaustive sense, even if at the assertive level the weakly exhaustive reading is the more salient.

We did not discuss in any systematic way what factors play a role in determining which of the two lexical rules are used. In the case of *know* and other factive verbs, our two rules generate, respectively, the strongly exhaustive reading and the weakly exhaustive reading. However, even though we are not sure whether a verb such as *know*, for instance, really licenses the weakly exhaustive reading, we insisted that the weakly exhaustive reading that is generated by a theory *à la Karttunen* (and ours) is in any case too weak. We briefly discussed how our lexical rule for the weakly exhaustive reading might be strengthened to deal with this problem, taking inspiration from recent work by Klinedinst & Rothschild. Finally, one type of reading we did not discuss at all in the paper is the so-called *mention-some* reading for embedded interrogatives, which is beyond the reach of our theory (see George 2011 for interesting discussions). While our proposal raises a number of unsolved problems, we hope to have clearly delineated what the major theoretical and empirical issues are for any future proposal aiming to improve on ours.

To conclude this paper, we would like to mention another foundational issue for research on embedded interrogatives. In order to reach a more complete understanding, we must not only provide a uniform theory of the meaning of the *V+interrogative* construction (which is the focus of our paper). We would also like to understand why some verbs, but not others, can take both a declarative and interrogative complement. For instance, why is it the case that one can say *John knows whether it is raining*, but not **John believes whether it is raining* ? If it turned out that the class of responsive predicates is more or less constant across languages, then an answer in terms of arbitrary syntactic selectional restrictions would clearly be insufficient. Rather, we would like to be able to predict which verbs and predicates that can take declarative complements can also take interrogative complements on the basis of some underlying semantic property. This is a topic for further research (see Egré 2008 for a tentative proposal).[30]

References

Baker C. L. (1968), *Indirect Questions in English*, Phd. Dissertation, University of Illinois at Urbana-Champaign.

Beck S. & Rullman H. (1999), "A Flexible Approach to Exhaustivity in Questions", *Natural Language Semantics* 7: 249-298.

Berman S. (1991), *The Semantics of Open Sentences*, Phd. Dissertation, University of Massachusetts.

Egré P. (2003), "Savoir, Croire et questions enchâssées", *Proceedings of Division of Linguistic Labor: The La Bretesche Workshop*, 2003, and Electronic Publications of *Philosophia Scientiae*.

Egré P. (2008), "Question-Embedding and Factivity", *Grazer Philosophische Studien* 77 (1), p. 85-125.

Fintel, K. von (2004), "Would you believe it? The King of France is Back! (Presuppositions and Truth-Value Intuitions)", in M. Reimer and A. Bezuidenhout eds., Descriptions and Beyond, Oxford University Press, pp. 315-341.

Fintel, K. & Heim, I. (2011), *Lecture Notes in Intensional Semantics*, Ms., MIT. Available at http://mit.edu/fintel/fintel-heim-intensional.pdf.

[30] Egré (2008)'s proposal adopts a strategy that is in fact symmetric to the one adopted here : the default meaning for all embedded questions is assumed to be veridical, and the existential meaning associated to non-veridical readings is derived by means of a semantic operation that takes place only if the embedded interrogative is the sister of a preposition.

Gajewski, J. (2005). *Neg-Raising: Polarity and Presupposition*, PhD Dissertation, MIT.

George B. (2011), *Question Embedding and the Semantics of Answers*, PhD Dissertation, UCLA.

Gettier, E. (1963). "Is Justified True Belief Knowledge", *Analysis*, p. 121-123.

Ginzburg J. (1995), "Resolving Questions : I & II", *Linguistics and Philosophy*, 18, pp. 459-527 & 567-609.

Groenendijk J. & Stokhof M. (1982), ``Semantic Analysis of Wh-Complements", *Linguistics and Philosophy*, 5, pp. 117-233.

Groenendijk J. & Stokhof M. (1990), ``Partitioning Logical Space", Course notes, ESSLLI 1990.

Groenendijk J. & Stokhof M. (1993), "Interrogatives and Adverbs of Quantification", in: K. Bimbo & A. Mate (eds), *Proceedings of the 4th Symposium on Logic and Language*.

Guerzoni, E. (2007), Weak Exhaustivity and Whether: A Pragmatic Approach, in T. Friedman & M. Gibson (eds.), *Proceedings of SALT*, Cornell, p. 112-129.

Guerzoni E. & Sharvit Y. (2007), A question of strength: on NPIs in interrogative clauses, *Linguistics and Philosophy* 30 (3): 361-391.

Heim I. (1994), "Interrogative Semantics and Karttunen's semantics for *know*", in R. Buchalla & A. Mittwoch, *IATL 1*, Hebrew University of Jerusalem, p. 128-144.

Heim I. & Kratzer, A. (1998), *Semantics in Generative Grammar*, Oxford: Blackwell.

Higginbotham J. (1996). "The Semantics of Questions", in S. Lappin (ed.), the Handbook of Contemporary Semantic Theory, Malden: MA, Blackwell.

Karttunen L. (1977), "Syntax and Semantics of Questions", *Linguistics and Philosophy*, 1: 3-44.

Kiparsky P. & Kiparsky C. (1970), "Facts", In M. Bierwisch and K.E. Heidolph (eds), *Progress in Linguistics*, pp. 143-73, The Hague: Mouton, 1970.

Lahiri U. (2002), *Questions and Answers in Embedded Contexts*, Oxford Studies in Theoretical Linguistics, Oxford.

Lewis D. (1982), " 'Whether' Report", in T. Pauli & al. (eds), *Philosophical Essays Dedicated to L. Aqvist on his 50th Birthday*, repr. in D. Lewis, *Papers in Philosophical Logic*, chap. 3, Cambridge Studies in Philosophy.

Rothschild, D. & Klinedinst, N. (2011), Exhaustivity in questions with non-factives, *Semantics and Pragmatics* 4.

Sharvit Y. (2002), "Embedded Questions and `De Dicto' Readings", *Natural Language Semantics*, 10: 97-123.

Schlenker P. (2006), “Transparency: An Incremental Theory of Presupposition Projection”, preliminary draft (18 pages), UCLA & Institut Jean-Nicod.

Schlenker, P. (2010), “Local Contexts and Local Meanings", *Philosophical Studies*.

Spector, B. (2005), Exhaustive interpretations: What to say and what not to say. Unpublished paper presented at the LSA workshop on Context and Content, Cambridge, MA.

Spector, B. (2006), *Aspects de la pragmatique des opérateurs logiques*, Thèse de l'Université Paris 7.

Spector, B. & Egré, P. (2007), Embedded Questions Revisited: *an* answer, not necessarily *the* answer. Handout, MIT linglunch, 2007, Nov. 8, available at http://lumiere.ens.fr/~bspector/Webpage/handout_mit_Egre&SpectorFinal.pdf