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

Benjamin SPECTOR & Paul EGRÉ
Département d'Études Cognitives
Institut Jean Nicod (ENS-EHSS-CNRS)
École Normale Supérieure – PSL Research University

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 responsive 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 *intermediate exhaustive readings* (cf. Klinedinst & Rothschild 2011) of embedded questions.

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).[1] In this paper, we are

---

* A version of the research reported here was 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. Our proposal has evolved very significantly since then.
We thank the editors of this special issue, Yacin Hamami and Floris Roelofsen, for their encouragement to submit our work, and for many helpful suggestions. We are grateful to the two reviewers of this paper for detailed and valuable comments, and in particular to Jeroen Groenendijk, whose extremely detailed and perspicuous comments contributed very significantly to the final version of our proposal. Special thanks go to Marta Abrusan, Emmanuel Chemla, Alexandre Cremers, Danny Fox, Benjamin George, Elena Guerzoni, Nathan Klinedinst, Daniel Rothschild, Savas Tsohatzidis, and to audiences at MIT (LingLunch 2007), Paris (JSM 2008), Amsterdam (2009), UCLA (2009), and the University of Maryland (2014). We also thank Melanie Bervoets, Heather Burnett, Nat Hansen and David Ripley for native speakers' judgments in English. The research leading to these results has received support from the Agence Nationale de la Recherche (grants ANR-10-LABX-0087 IEC, ANR-10-IDEX-0001-02 PSL and ANR-14-CE30-0010-01 TriLogMean)..

[1] We don't know yet of a comprehensive account of embedded questions within the framework of *Inquisitive Semantics* (see Ciardelli et al. 2013 for a survey). However, Ciardelli & Roelofsen (this issue) provide a treatment of embedding of declaratives and interrogatives under an epistemic operator representing knowledge, as well as under *wonder*, within the framework of *Inquisitive Semantics*. See also Theiler (2014) and Roelofsen et al. (2014).

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,[2] no currently available theory is able to achieve what would seem like the main desideratum, namely to characterize in full generality, *without specific, idiosyncratic lexical stipulations*, the semantic relationship between the *V+declarative* construction and the *V+interrogative* construction, for any *V* that can embed both declarative and interrogative sentences. We will sketch such a theory. Our goal is relatively modest: we want our theory to *formalize* some important empirical generalizations. We hope that the final, correct theory of interrogatives and interrogative embedding will have to derive these generalizations as well, even if it turns out to have a very different shape than ours.

Let us be more specific. Consider the two following pairs:

(1) a. Mary knows that it is raining.
b. Mary knows whether it is raining.

(2) a. Mary told Paul that it is raining.
b. Mary told Paul whether it is raining.

In both cases, a certain attitude verb (*know*, *tell*) embeds either a declarative sentence or an interrogative sentence. We would like our theory to be able to predict the truth-conditions of these four sentences on the basis of a) the meaning of the attitude verb, b) the meaning of the declarative clause ‘it is raining’, and c) the meaning of the interrogative clause ‘whether it is raining’. Moreover, we want the theory to do this by means of general compositional rules that do not refer to the identity of the relevant attitude verbs, i.e. they should predict the meaning of ‘*V+interrogative*’ and ‘*V+declarative*’ for any *V*, in terms of the lexical meaning of *V* and the meaning of its complement. Now, some conceivable ways of doing this would be devoid of any interest. For instance, we may design a framework in which we can associate to every such *V* an ordered set of two functions, one corresponding to the case where it embeds a declarative and one to the case where it embeds an interrogative. Consider for instance the following fictitious lexical entry, where we use ‘?’ as the type for interrogatives, given that at this point we have not said anything about the meaning of interrogatives:[3]

(3) $[\![\text{shknow}]\!] = < \lambda p_{<s,t>}. \lambda x.x \text{ knows } p, \lambda Q_{?}.\lambda x.x \text{ wonders } Q>$

According to this lexical entry, *John shknows that it is raining* is equivalent to *John knows that it is raining*, and *John shknows whether it is raining* is equivalent to *John wonders whether it is raining*. Obviously, this would be no better than simply stating two independent lexical entries for each case. Rather, what we are after is this. A verb like *know* denotes a propositional attitude. We want to characterize the meaning of sentences of the form *X knows that S* and *X knows Q* in terms of this propositional attitude and the meaning of *S* and *Q*, and we want to do this in a way that is generalizable to other attitude verbs. The hope is that once this is done, were we to cash out the meaning of these attitude verbs in terms of pairs of functions, as in (3) (but why would we?), certain conceivable pairs (such as, presumably, (3) itself) would be simply impossible given the theory. In other words, our empirical goal is to impose substantial constraints on the possible meanings of pairs of sentences of the form < *X shknows that S*, *X shknows whether S* >.

[2] The exception is Ben George (2011), who partly builds on the ideas we present here. Though we do not discuss George’s specific proposal, many aspects of the second half of our paper are related to George’s work.
[3] We use here semi-formal lexical entries. This is meant only for illustration.

Now, various possible strategies are conceivable to achieve this goal. One strategy might for instance adopt the 'pair' approach we have just outlined, but add specific constraints to rule out certain pairs. Another conceivable approach would consist in stating a lexical entry corresponding to the case where the attitude verb embeds an interrogative, and to give a general rule determining the meaning that results from combining the attitude verb with a declarative complement. Conversely, we might assume that the basic meaning of the relevant attitude verbs corresponds to a *propositional* attitude and give a general rule determining the meaning that results from combining the attitude verb with an interrogative complement. In one version of this approach, each attitude verb would be associated with a pair of meanings, as in (3), but each member of the pair would be derived by a *general rule* from a propositional attitude serving as the 'core' meaning of the verb. Another variant would consist in directly deriving the meaning of *V+interrogative* in terms of the meaning of *V+declarative*. Yet another possibility would be to deny that *declaratives* and *interrogatives* have different semantic types, and to give only one single lexical entry for each verb and a simple compositional rule that would work both when the verb embeds a declarative and an interrogative complement. Such a strategy is particularly natural within the new framework known as *Inquisitive Semantics* (Ciardelli 2009, Ciardelli and Roelfosen 2009, Mascarenhas 2009, Groenendijk 2009, Groenendijk & Roelofsen 2009, Ciardelli et al. 2013, Roelofsen 2013, among others), in which an enriched notion of meaning is used with the result that declaratives and interrogatives denote the same type of object. In such a framework, one might hope that it might be impossible to define a lexical entry for *shknow* whose end result would be empirically equivalent to (3). Possibly one might need to use meaning postulates to narrow down the possible lexical entries.[4]

In this paper, we do not develop and compare these different strategies.[5] We will simply develop one such strategy, namely the one that consists in deriving the meaning of the *V+interrogative* construction in terms of the meaning of the *V+declarative* construction. In so doing, we are in the footsteps of some of the classical works on interrogative semantics (Karttunen 1977, Groenendijk & Stokhof 1982). We should note that such a strategy makes a pretty strong assumption regarding the semantic relationship between the two constructions. Namely, it implies that for all predicates P that take both declarative and interrogative complements (*responsive predicates* in Lahiri's 2002 typology), the meaning of the *P-interrogative clause* construction is predictable from the meaning of the *P-declarative clause* construction. As we shall see, the current state of the literature does not provide such a characterization.

Our account 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

[4] We are indebted to J. Groenendijk for this comment.
[5] But see George (2011) p. 142-169 for a discussion of such issues.

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. Such a view was challenged by Tsohatzidis (1993, 1997), who argued that embedded questions after *tell* do not always have a veridical reading. In support of Tsohatzidis’ view, 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 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 favor *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 responsive predicates, we can give a plausible account of weakly exhaustive readings. We argue 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. Sections VI and VII further refine our proposal. Section VI presents a first refinement intended to account for verbs like *discover*, which trigger a negative presupposition (*discover S* presupposes *not knowing S in the past*), based on ideas originally developed (independently) by Jeroen Groenendijk (p.c.) and Danny Fox (p.c. and classnotes, Fox 2013). Section VII discusses a problem for the usual characterization of weakly exhaustive readings, and provides a slightly stronger characterization that makes them equivalent to what we call, following Klinedinst & Rothschild (2011), *intermediate readings*. The final form of our account appears in a single meaning postulate stated under (130). We also discuss possible modifications of our final proposal in a separate appendix. Though we consider sections VI to VII to be essential to our theory, readers not concerned with those refinements may already find a fair approximation of our account at the end of section V (section V.6, statement (116)).

The resulting theory, however, will not be as constrained as one could hope. Our system will be able to generate both *strongly exhaustive* readings and *weakly/intermediate* readings. However, there are arguments in the literature that some predicates license only one type of reading, and there is also massive disagreement about the facts. For instance, Groenendijk and Stokhof (1982) and George (2011) argue that only strongly exhaustive readings exist, while Guerzoni & Sharvit (2007) claim that emotive factive verbs only license weakly exhaustive readings but that all other relevant predicates license both readings. Finally, Klinedinst & Rothschild (2011) suggest that all these readings are always available in principle. Recently, Cremers & Chemla (2014) provided experimental evidence that the intermediate reading is available for a verb such as *know*. While we will discuss some of the relevant empirical issues, our proposal will be in itself compatible with all these views, if conceived as a theory that significantly constrains but does not fully determine the space of possible meanings for the *V+interrogative* construction.

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):[6]

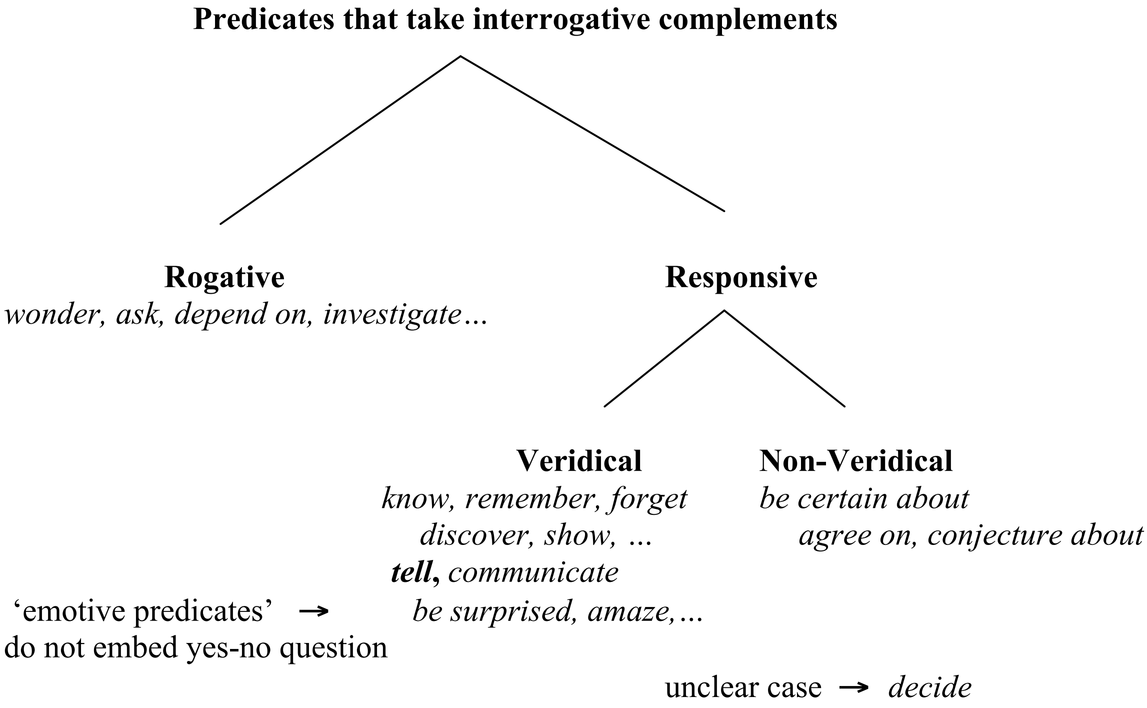

Rogative predicates are predicates that take interrogative complements and do not license declarative complements. We will not be concerned with such predicates.

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

(i) Syntactic property: they take both declarative and interrogative complements
(ii) Semantic property: 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

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

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

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

(6) 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

(7) "John knows whether it is raining" is true iff John knows p, where p is the correct answer to "is it raining ?"
(8) "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 Groenendijk & Stokhof's sense)

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

Illustration:

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

(10) 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

(11) Sue is certain about who came
→ can be true even if Sue's belief about who came is false

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:

(12) 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,[7] with respect to 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 verbs such as *tell*, *announce, predict.* As observed by many authors (Karttunen 1977, Groenendijk & Stokhof 1982, Lewis 1982, Berman 1991, Higginbotham 1996, Lahiri 2002), while (13) below does not entail that what Jack said to Mary is the truth (i.e. does not entail that Peter is the culprit), (14) 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*.

(13) Jack told Mary that Peter is the culprit.
(14) 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 (12)a:

(15) Veridical-responsive predicates are exactly those responsive predicates that are factive or veridical with respect to their declarative complements.[8]

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 (12)b. Before offering such an answer, though, we want to show that our generalization in (15) can be defended against the apparent counterexamples mentioned above, i.e. the behavior of communication verbs.

---

[7]Let us 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 if 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)

[8] 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. 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 (15).

II.1 Communication verbs are not systematically veridical-responsive

As observed above, a sentence like (16) 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:

(16) 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"):

(17) "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*. Tsohatzidis (1993) gave the following pair as an example of a nonveridical use of *tell+question,* and contrasted it with the veridical reading of questions under verbs such as *reveal* or *remind*. Tsohatzidis points out that (18) is not contradictory, and that (19) is not redundant:

(18) Old John told us whom he saw in the fog, but it turned out that he was mistaken (the person he saw was Mr. Smith, not Mr. Brown)
(19) Old John told us whom he saw in the fog, and it turned out that he was not mistaken (the person he saw was indeed Mr. Brown)

Further examples can be given in support of Tsohatzidis' observations:

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

(22) I believe that Jack told you which students passed; but I don't think he got it right.

(23) I believe 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:

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

While (24) strongly suggests that Jack made a correct prediction, (25) and (26) are felicitous, but would 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 (16) and (24). 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 2007).

Consider the following sentences:

(27) Sue told someone that she is pregnant.
(28) Sue didn't tell anyone that she is pregnant.
(29) Did Sue tell anyone that she is pregnant?

From (27), one easily infers that Sue is in fact pregnant. This is of course no argument for the view that (27) 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 (27) 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, (28) and (29) 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 actually be 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.[9]

---

[9] 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 (29). According to some informants if a dative argument is added for

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:

(30) Sue told Jack that Fred is the culprit.
(31) Sue didn't tell Jack that Fred is the culprit.
(32) Did Sue tell Jack that Fred is the culprit?

While (30) may or may not trigger, depending on context, the inference that Fred is the culprit, (31) and (32) 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 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.

(33) The Wait a Minute Test ('WMT' 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 WM test to (30), (31) and (32):

(34) - Sue told Jack that Fred is the culprit
- Hey wait a minute! I didn't know that Fred is the culprit
(35) - Sue didn't tell Jack that Fred is the culprit
- Hey wait a minute! I didn't know that Fred is the culprit
(36) - 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 S* does not behave the way *tell someone that S* does:

(37) Sue said that Fred is the culprit
(38) Sue didn't say that Fred is the culprit
(39) 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, (37) suggest that Fred is in fact the culprit, but (38), 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*.

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 still not to the same extent as (29).

Whatever the source of this complex behavior is, we may safely conclude that *tell someone that S* 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).[10]

Taking into account the facts pointed out in section II.1, we have shown that *tell* has the following properties:

(40) 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 (15), repeated as (41):

(41) 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.2.2. Other communication verbs: *Announce*, *Predict*

Consider now:

(42) Mary announced that she is pregnant.
(43) Mary didn't announce that she is pregnant.
(44) Did Mary announce that she is pregnant?

(42) 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 (43) and (44) 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:

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

---

[10] An anonymous reviewer suggests that the verb 'suspect' likewise may have both a factive and a non-factive use, but is not in any clear sense a 'communication verb'. The same appears to hold of 'guess' in English (see Egré 2008). 'Guess', like 'suspect', is not a communication verb ; unlike 'guess', however, 'suspect' does not appear to license embedded questions.

There seems to be a reading for (45) which licenses the inference that all of these ten girls are pregnant, much like "None of these ten girls knows that she is pregnant".

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

(47) - 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:

(48) Mary predicted that she would be pregnant.
(49) Mary didn't predict that she would be pregnant.
(50) Did Mary predict that she would be pregnant?
(51) None of these ten girls had predicted that she would be pregnant.

While (48), depending on context, may or may not trigger the inference that Mary got pregnant, this inference is easily drawn, out of the blue, from (49) and (50). And (51) strongly suggests that the speaker takes the ten girls he is referring to to be actually pregnant (or to have been).

Again, the WMT confirms that we are dealing here with presupposition-like inferences: both (49) and (50) license an objection of the following form:

(52) 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:

(53) 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. 3. *Predict* and *Guess* in French

In support of the generalization stated in (41), we should point out an interesting French minimal pair, made up of the two verbs *predire* ('predict') and *deviner* (approximately factive 'guess'). 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,[11] 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

[11] 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*.

deviner does not.[12] Now, it turns out that both *prédire* and *deviner* can take interrogative complements. What (41) implies is that *deviner*, but not *prédire*, is veridical-responsive. This prediction is clearly borne out :

(54) (?) 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.*

(55) # Marie a deviné qui viendrait à la fête, mais elle s'est trompée.
Marie guessed who would attend the party, but she got it wrong.

Although (54) 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 (54) and (55). And the following judgments, which make the same general point point, are accepted by all our informants:

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

(57) 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.4. *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.).

(58) 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.

(59) 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'.

---

[12] 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.

→ The Eiffel tower will collapse.

Now, both *mond* and *elmond* embed interrogative complements. As we expect given our generalization in (41), *mond* is not veridical-responsive, but *elmond* is. The relevant judgments are as follows. (60) shows that *elmond* is veridical responsive.

(60) 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.

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).

(61) 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.5. 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 (41), 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 (41). 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:

(62) 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 write Q(w) for the complete answer to Q in w. We thus view a question Q as denoting a function from worlds to propositions – type <s, <s,t>>. It will also be convenient to introduce a notation standing for the set of *potential complete answers to Q*. We denote this set as *pot(Q)*, defined as follows:

(63) $\text{pot}(Q) = \{p : \exists w\ (p = Q(w))\}$

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 (P_{int} is thus of type <<s,<s,t>>, <e,t>>).[14] Our proposal is captured by the *general* meaning postulate in (64)a, which is equivalent to (64)b:[15]

(64) a. $[[P_{int}]]^w = \lambda Q.\lambda x.\ \exists p \in \text{pot}(Q)\ [[P_{decl}]]^w(p)(x) = 1$
b. $[[P_{int}]]^w = \lambda Q.\lambda x.\ \exists w'\ [[P_{decl}]]^w(Q(w'))(x) = 1$

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, (64) 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:

(65) $[[P_{G\&S\text{-}int}]]^w = \lambda Q.\lambda x.\ [[P_{decl}]]^w(Q(w))(x) = 1$

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 (41), since it leads us to expect that every responsive predicate is veridical-

[14] 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. 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. It should be 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, as in Egré (2008), which is based on similar ideas as this paper (however, it is not entirely clear how the facts we discuss in section IV about presupposition projection, and our proposal developed in the final sections regarding the ambiguity between weakly/intermediate and strongly exhaustive readings can be straightforwardly accommodated in these terms). Another possible approach, discussed by Jeroen Groenendijk in an extensive review of the present paper, would consist in moving to a framework where declarative and interrogative sentences have the same semantic type. We do not discuss these issues in this paper. 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, with the hope that the gist of our proposal can be captured in such alternative frameworks.

[15] 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. Later in this paper, as we will introduce several notions of complete answers, our rules will take a slightly different form

responsive.[16] 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.[17] 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 (extensional 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 (65). We argued, however, in agreement with Tsohatzidis (1993, 1997), that verbs such as *tell* are only apparent exceptions to the generalization (41), 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 (64) 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 clause is embedded. In particular, Groenendijk & Stokhof's (1982) definition of a 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. Various notions of complete answer - weak and strong exhaustivity

In the literature on the semantics of interrogatives, two distinct notions of what counts as a *complete answer* to a question are usually discussed. 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 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). In this section, we introduce various notions of complete answer, on the basis of a more primitive one, which we take to be the 'basic' denotation of questions. Our reasons for doing so are expository. We then show that the most straightforward implementation of the proposal we have sketched in the previous section requires that we use the 'strongly exhaustive reading' as the relevant notion of complete answer. However, later in this paper, we will argue that once we take into account the presuppositional behavior of question-embedding predicates, it will be both possible and desirable to propose *several* meaning postulates, based on different notions of complete answers.

Let us first introduce various possible notions of complete answer that we will discuss in this paper. Note, first, that they are not intended to make any difference for *polar questions*. We represent polar questions (such as *Is it raining?*) as $?\phi$, where ϕ is the declarative sentence on the basis of which the polar question is formed. We adopt a completely standard semantics, according to which the complete answer to $?\phi$ is the proposition expressed by ϕ if ϕ is true, its negation otherwise. In other words, in a world w, the complete answer to $?\phi$ is the set of worlds v in which ϕ has the same truth-value as in w. Taking the denotation of a polar question in a given world to be the complete answer to the question in that world, we thus have the following rule (see Groenendijk and Stokhof 1982):

---

16 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.

17 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).

(66) $[[?\phi]]^w = \lambda v. [[\phi]]^v = [[\phi]]^w$

It is for constituent questions that we will make use of several different notions. We only deal here (for simplicity) with constituent questions containing just one *wh*-word (as opposed to multiple constituent questions), and we adopt for such questions a very simplified representation. We assume that the *wh*-word always combines with a constituent of type <e,t>, binds a variable of type *e*, and we treat the restrictor of the *wh*-word as a sortal restriction on the variable. As a result, we represent a constituent question such as *Which students left* as in (67)b. The general form of a constituent question is thus given in (67)c:

(67) a. Which students left?
b. $?x_{student}$ LEFT(x)
c. $?x_R$ P(x)

In practice, we will often ignore the restrictor, which will play no crucial role in the subsequent discussions, and will then simply represent constituent questions as ?x P(x).[18]

Regarding constituent questions, we can start with a first notion of *complete answer*, which is close but not identical to Karttunen's notion of complete answer (Karttunen 1977), and which will serve as the basis for more complex notions. This most basic notion of complete answer, which we take to be the *basic denotation of a constituent question in a world of evaluation w*, is defined as in (68)a, which is equivalent to (68)b - D^w is the domain of individuals in *w*.

(68) a. $[[?x_R\ P(x)]]^{w,g} = \lambda v. ([[R]]^w \cap [[P]]^w) \subseteq ([[R]]^w \cap [[P]]^v)$
b. $[[?x_R\ P(x)]]^{w,g} =$
$\lambda v. \forall d \in D^w [([[R(x)]]^{w,g(x/d)}=1 \text{ and } [[P(x)]]^{w, g(x/d)} = 1) \rightarrow (d \in D^v \text{ and } [[P(x)]]^{v, g(x/d)} = 1)]$

Ignoring the restrictor, we get:

(69) a. $[[?x\ P(x)]]^{w,g} = \lambda v.[[P]]^w \subseteq [[P]]^v$
b. $[[?x\ P(x)]]^{w,g} = \lambda v.\forall d \in D^w [([[P(x)]]^{w, g(x/d)} = 1) \rightarrow (d \in D^v \ \& \ [[P(x)]]^{v,g(x/d)} =1)]$

In words: the denotation of a question of the form *?x P(x)* in a world *w* is the proposition that states, for every individual *d* who has the property *P* in world *w*, that *d* has property *P*. Taking into account restrictors, $?x_R$ *P(x)* denotes in a world *w* the proposition that states, for every individual *d* that has the properties expressed by *R and P* in *w*, that *d* has property *P*. For instance, the question *Which linguists left?*, evaluated in world *w*, denotes the conjunction of all propositions of the form *x left*, where *x* is a linguist who left in *w* (note that such a proposition does not itself entail that *x* is linguist). This denotation for questions corresponds to what can be thought of as the most basic notion of a complete answer, which we can call the *mention-all answer*. When referring to this notion of complete answer, we will use the phrase *basic complete answer*. A proposition is the actual basic complete answer to '?x P(x)' in a world *w* if and only if it says, for every x that has the property P in *w*, that x is P. Note that in case P has an empty extension in *w*, then the denotation of the question in *w* is the tautology (this is so because the condition $[[P]]^w \subseteq [[P]]^v$ is satisfied by every world *v* when $[[P]]^w = \emptyset$).

[18] In particular, we do not discuss in this paper the ambiguity between the *de re* and *de dicto* readings of *wh*-phrases in embedded questions. We are only considering contexts where the denotation of the restrictor is common knowledge, in which case the two readings collapse.

On this basis, one can define a stronger notion of complete answer, namely G&S's notion, which we will call the *strongly exhaustive complete answer*, or simply the *strongly exhaustive answer*. Consider an operator, noted exh_Q, which, when applied to a proposition *p*, returns the proposition that states that *p* is the basic complete answer to *Q* [exh_Q is used in the metalanguage, not in the object language]. For instance, if the question is *who came?*, then exh_Q(⟦*Mary and John came*⟧) is the proposition that states that *Mary and John came* is the basic complete answer to *who came?* – in other words, the proposition that states that Mary and John came and nobody else did.

(70) $exh_Q(p) = \lambda w.(Q(w) = p)$

Now, the notion of complete answer adopted by partition semantics corresponds exactly to this strengthening procedure. According to this notion, the complete answer to *who came?* is the proposition that states, for every *x* who came, that *x* came, and also states that nobody else came. So the G&S denotation of a question Q in a world *w*, noted GS(Q)(w), can be defined in terms Q(w) as in (71)a, which is equivalent to (71)b

(71) a. $GS(Q)(w) = exh_Q(Q(w))$
b. $GS(Q)(w) = \lambda v.(Q(v) = Q(w))$

Thus, *GS(Q)(w)* happens to denote the set of worlds in which the basic complete answer is the same as in *w*. And GS(Q) (putting aside irrelevant type-theoretical issues) expresses an equivalence relation over the set of worlds. Let us now consider what happens to a question *?xP(x)* in a world *w* where *P* has an empty denotation. Note that $[\![?x\ P(x)]\!]^w$ denotes the tautology (noted T) *if and only if* the extension of *P* is empty in *w*. So we have:

(72) $GS([\![?x\ P(x)]\!])(w) =$
$exh_{?x\ P(x)}(T) =$
$\lambda v.([\![?x\ P(x)]\!]^v = T) =$
$\lambda v.\ [\![P]\!]^v = \emptyset$ [i.e., the proposition that states that P has an empty extension]

Now, Karttunen's notion of complete answer is very close to our basic denotation for *wh*-questions, with one important difference. In the case where nobody came, K. defines the complete answer to 'Who came?' as being the proposition that states that nobody came. Thus, in this case, K's complete answer is the same as G&S's complete answer. If, however, at least one person came, K's complete answer is simply the proposition that states, of every person who came, that this person came. Karttunen's notion of complete answer can thus be defined in terms of our primitive notion, by means of the following operator *Kart*:

(73) $Kart_Q(p) = p$ if p is not T (the tautology), and $Kart_Q(p) = \lambda v.(Q(v) = T)$ if p = T.

Applied to any non-tautological proposition, $Kart_Q$ is just the identity function. Applied to the tautology, $Kart_Q$ returns the proposition that states that the complete answer to Q is the tautology, i.e., if Q is `Who came', the proposition stating that nobody came.

Now we can define the Karttunen denotation of a question as follows:

(74) a. $K(Q)(w) = Kart_Q(Q(w))$
b. $K(Q)(w) = \lambda v.(Q(v) = 1$ and Q(v) is not the tautology) or (Q(v) is the tautology and $Q(w)=Q(v)$

The motivation for Karttunen's notion of complete answer is the following. According to Karttunen, a sentence such as *John knows who came* is true as soon as John knows, for every person who came, that this person came (i.e. John doesn't need to know that *only* the people who actually came came). This reading, whose existence is highly debated (see section V), is called the *weakly exhaustive reading*. However, one thing that seems clear is that if in fact nobody came and John has no idea whatsoever about who came, one doesn't want *John knows who came* to be true by virtue of the fact that John knows that the tautology is true. Karttunen thus defined the complete answer as being identical to our basic notion of complete answer, except in the case where the basic complete answer would turn out to be the tautology – in which case K defines the complete answer as would G&S.[19]

Note that the GS operator when applied to the basic denotation of a question yields the same result as when it is applied to the Karttunen-denotation of a question. That is:

(75) For any world w and any question Q, GS(K(Q))(w) = GS(Q)(w)

Finally, let us briefly come back to polar questions. Given a polar question $?\phi$, and the rule in (66), we have the following identities:

(76) a. $GS([\![?\phi]\!]) = K([\![?\phi]\!]) = [\![?\phi]\!]$
b. For any proposition p in $\{[\![\phi]\!], [\![\phi]\!]\}$, $exh_{?\phi}(p) = Kart_{?\phi}(p) = p$[20]

From this it follows that the different notions of complete answer we have introduced are all equivalent for polar questions.

III. 2. The need to use strongly exhaustive answers.

We want *X V Q* to mean `X is in the relation V to the complete answer to Q', but we need to know which notion of complete answer we want to use. To make things simple, we will define everything in terms of the basic notion of answer, and will not directly refer to K(Q) and GS(Q) in formulating various rules. Rather, we will use the operators *exh* and *Kart* and apply them to potential basic complete answers.

Let us now use the definition of *pot(Q)* as {p: for some w, p = Q(w)}, where Q(w) is the denotation of Q in w according to the most basic notion (*mention-all*). That is *pot(Q)* denotes the set of potential basic complete answers**.**

We need to choose between the following possibilities.

(77) a. $[\![P_{int}]\!]^w = \lambda Q.\lambda x.\ \exists p \in pot(Q)\ [\![P_{decl}]\!]^w(p)(x) = 1$
b. $[\![P_{int}]\!]^w = \lambda Q.\lambda x.\ \exists p \in pot(Q)\ [\![P_{decl}]\!]^w(Kart_Q(p))(x) = 1$

[19] There is another notion of complete answer that one might consider, according to which the complete answer is not defined when the basic complete answer is a tautology. This would correspond to a claim that *wh*-questions carry an existence presupposition, an issue that has been much discussed (see, e.g., Dayal 1996). We do not want to endorse this claim and do not discuss this option further in this paper – except in footnote 35 and 39.

[20] Proof of (76)b. First, unless ϕ is tautologous, $Kart_{?\phi}([\![\phi]\!]) = [\![\phi]\!]$ by definition.

If S is tautologous then $Kart_{?\phi}([\![\phi]\!]) = Kart_{?\phi}(T) = \lambda v(T=([\![?\phi]\!])(v)) = T$.

We also have: $(exh_{?\phi}([\![\phi]\!])(u) = 1 \iff ([\![?\phi]\!](u) = [\![\phi]\!] \iff (\lambda v.\ [\![\phi]\!](v) = [\![\phi]\!] \iff [\![\phi]\!]^u = 1$. Symmetrically with $\neg\phi$. (76)a. follows straightforwardly.

c. $[\![P_{int}]\!]^w = \lambda Q.\lambda x.\ \exists p \in pot(Q)\ [\![P_{decl}]\!]^w(exh_Q(p))(x) = 1$

Now, (77)a and (77)b are both inadequate to capture the interpretation of embedded *wh*-questions. This is so because both rules would validate the following inference pattern:

(78) <u>John knows that Mary left</u>
John knows who left

This is a bad result. In fact, using (77)b, the predicted meaning for *John knows who left* is *Either nobody left and John knows this, or there is someone who left such that John knows that this person left.* Using (77)a would be even worse: the predicted meaning would then simply be that the sentence should be equivalent to *John knows that the tautology is true*, which we may argue is itself a tautology (this is so because the tautology is a potential complete answer in the basic sense, and so it would be sufficient for John to know the tautology for the sentence to be true). The general point here is that using either of these two notions of complete answer, together with our general perspective on question embedding, does not in fact generate the weakly exhaustive reading as defined by Karttunen, but something much weaker. This is so because by quantifying existentially over potential basic answers, we lose the condition that the answer that has to be known has to be the *actual complete answer*. It now only has to be a true *potential basic complete answer* (true because *know* is factive). But a potential complete answer in the basic sense can be true without being the actual complete answer. For instance, if both John and Mary left, *John left* is a true potential basic complete answer to *Who left?*, but is not the actual basic complete answer.

On the other hand, on the basis of (77)c, we derive the same result as Groenendijk & Stokhof for *know*. This is so because the *strong* notion of complete answer *(i.e* $exh_Q(p)$, where p is a complete answer in the weak sense), unlike the basic notion, is such that a strongly exhaustive answer cannot be true unless it is the *actual* strongly exhaustive complete answer.[21] Another way of saying this is that strongly exhaustive answers are mutually exclusive. For instance, if the people who left are John and Mary, then there is only one true strongly exhaustive answer, namely the proposition that John and Mary left and nobody else did. Now, according to (77)c, *John knows who left* is true if for some potential answer (in the basic sense) p, John is in the relation know to $exh_Q(p)$. And for John to be in the relation *know* to some proposition $exh_Q(p)$, $exh_Q(p)$ has to be true, because *know* is veridical. But in a given world, there is only one potential answer (in the basic sense) p such that $exh_Q(p)$ is true: $exh_Q(p)$ is true in world u if and only if $p = Q(u)$. So if there is a p such that *John knows* $exh_Q(p)$, $exh_Q(p)$ has to be the unique *true* strongly exhaustive answer to Q. It thus follows that *John knows Q* ends up equivalent to *John knows the truth of what is in fact the complete answer to Q in the G&S' sense*. We have illustrated here a more general instance of the following fact.

(79) If *P* is a *veridical* predicate, then the following two statements are equivalent, for any question Q and any world *w*:
a. $\exists p \in pot(Q)\ ([\![P_{decl}]\!]^w(exh_Q(p))(x) = 1$
b. $[\![P_{decl}]\!]^w(GS(Q)(w))(x) = 1$

Now, while (79)a is our general scheme for question-embedding (based on strongly exhaustive answers), (79)b is equivalent to G&S's proposal. Our point is that even though

[21] Namely, if $exh_Q(p)(w) = 1$, then $p = Q(w)$ and $exh_Q(p) = exh_Q(Q(w))$. <u>Proof</u> : assume $exh_Q(p)(w) = 1$. Then by definition $\lambda v(Q(v) = p)(w) = 1$, i.e. $Q(w) = p$.

there is no general equivalence between the two approaches, the two approaches are equivalent when the embedding predicate is veridical (and in particular when it is factive).

At this point, it seems that we are forced to use the strong notion of complete answer (*strongly exhaustive complete answer*), since using a weaker notion gives rise to patently too weak truth-conditions. That is, because our lexical rule involves existential quantification over potential complete answers, we would clearly predict too weak truth-conditions if we used either the basic notion or Karttunen's notion of complete answer.

Our official meaning postulate for defining the meaning of P_{int} in terms of P_{decl} is thus, at this point, the following:

(80) $[[P_{int}]]^w = \lambda Q.\lambda x.\ \exists p \in pot(Q)\ [[P_{decl}]]^w(exh_Q(p))(x) = 1$

This does not capture the reading initially predicted by Karttunen, called *the weakly exhaustive reading*, which has been argued to exist as well (Beck & Rullman 1999, Guerzoni & Sharvit 2007, and, more recently Klinedinst & Rothschild 2011). In the case of *know*, this reading can be paraphrased as follows:

(81) *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 (81) 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 (81). Yet, like G&S, we predict the sentence to be false. So far our point is simply that our proposal cannot capture this reading, whether or not it really exists – because given the general shape of our proposal (existential quantification over potential complete answers), using either the basic notion of a complete answer or Karttunen's slightly stronger notion does not give rise to this reading, but to a much weaker one.

As we shall discuss later, this is problematic in light of data from systematic surveys that show, conclusively in our view, that a reading which is very close to the *weakly exhaustive readings* (namely the *intermediate reading*, see section VII) exists as well, at least for some responsive predicates (Klinedinst & Rothschild 2011, Cremers & Chemla 2014). Furthermore, *surprise*-type verbs (so-called emotive factive predicates) have been argued to *only* support weakly exhaustive readings (Guerzoni and Sharvit 2007). We offer a solution to this problem in section V, which we then refine in sections VII and VIII.

IV. Incorporating presuppositions

Suppose now that P_{decl} is a presupposition trigger, i.e. that a sentence of the form 'x P S', where x and S are not themselves presuppositional, has 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 (80). Clearly, for the condition '$[\![P_{decl}]\!]^w(exh_Q(p))(x) = 1$' to hold, $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ has to be defined in the first place. Now, let *P* be a factive predicate such as *know*. Assume for simplicity the following lexical entry for *know*$_{decl}$, according to which 'x knows S' *presupposes* that S is true and is otherwise equivalent to 'x believes S'.[22]

(82) $[\![know_{decl}]\!]^w = \lambda p.\lambda x\text{: } p(w) = 1.\ x \text{ believes } p \text{ in } w.$[23]

Given (82), the condition "$[\![know_{decl}]\!]^w(exh_Q(p))(x) = 1$" is equivalent to "$exh_Q(p)(w) = 1$ and x believes $exh_Q(p)$ in w". We thus have:

(83) $[\![know_{int}]\!]^w=\lambda Q.\lambda x.\exists p \in pot(Q)(exh_Q(p)(w)=1\ \&\ x \text{ believes } exh_Q(p) \text{ in } w)$

Now the condition '$exh_Q(p)(w) = 1$' simply means that $p = Q(w)$. So we end up with:

(84) $[\![know_{int}]\!]^w = \lambda Q.\lambda x.\ \exists p \in pot(Q)\ (p = Q(w) \text{ and } x \text{ believes } exh_Q(p) \text{ in } w)$

i.e.:

(85) $[\![know_{int}]\!]^w = \lambda Q.\lambda x.\ x \text{ believes } exh_Q(Q(w)) \text{ in } w$

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 (82).

IV.1. A presuppositional variant

So far (80) predicts that even when P_{decl} is a presupposition trigger, P_{int} will not be. Rather, what follows from (80) is that for some potential strongly exhaustive 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.

(86) $[\![P_{int}]\!]^w = \lambda Q.\lambda x\text{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(exh_Q(p))(x) \text{ is defined}).$
$\exists p\in pot(Q)\ ([\![P_{decl}]\!]^w(exh_Q(p))(x) = 1)$

In more informal terms, what (86) says is the following:

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

---

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

[23] We use Heim & Kratzer's 1998 notation for representing presuppositions in our lexical entries. Namely, a lexical entry of the form $[\![X]\!]^w = \lambda A. \ldots . \lambda Z\text{: } \phi(A,\ldots,Z).\psi(A,\ldots,Z)$, 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 $\phi(A,\ldots,Z)$ is true, and, when defined, returns the value 1 if and only if $\psi(A,\ldots,Z)$ holds. In other words, the presuppositions triggered by an expression are encoded in the part of the formula standing between the column and the period. We will often call the right-hand side of such expressions the 'assertive part' of an expression, but this is in fact improper, as discussed in footnote 32.

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 (87)a to *know* predicts the following presupposition for 'x knows Q':

(88) 'x knows Q' presupposes that for some potential strongly exhaustive answer S to Q, the presupposition of 'x knows S' is true.

This is equivalent to:

(89) 'x knows Q' presupposes that for some potential strongly exhaustive answer S to Q, S is true.

But note that the presupposition predicted by (89) is a tautology: it simply states that Q has a true strongly exhaustive answer, which is necessarily the case given the kind of semantics for questions we are assuming.

Yet (86) and (80) 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 (86).

IV. 2. *Agree that/Agree on*

In this section, we will be concerned with the presuppositions and truth-conditional content of sentences in which a question is embedded under *agree on* – a type of case discussed at length in Lahiri (2002), and in Beck & Rullman (1999) as well as Sharvit (2002).

We first focus on the construction *X agrees with Y on Q* (which is not discussed by Beck & Rullman 1999 or Sharvit 2002), and then on its reciprocal variant *X and Y agree on Q*.

To determine the predictions of our proposal in (86), we need to know what presuppositions *agree* yields when it takes a declarative complement, as in (90):

(90) Jack agrees with Sue that it is raining

We follow Lahiri in assuming that (90) presupposes that Sue believes that it is raining and asserts that Jack shares this belief. That is, we posit the following lexical entry for *agree*:

(91) $[\![\text{agree}_{\text{decl}}]\!]^{w} = \lambda p.\lambda y.\lambda x{:}(Dox_y(w) \subseteq p).\ Dox_x(w) \subseteq p$
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).

One might wonder whether (90) does not also presuppose that Jack is *opinionated* with respect to whether it is raining (i.e. either believes that it is raining or believes that it is not raining). This might explain why the negation of (90) (*Jack doesn't agree with Sue that it is raining*) does not only suggest that Jack does not share Mary's belief that it's raining, but tends to imply that Jack believes that it's not raining. This is reminiscent of the *neg-raising* reading of *believe* (*X doesn't believe S* suggests that X believes that S is false), and neg-raising has been treated in presuppositional terms (Gajewski 2005). However, other approaches do not assume that neg-raising should be captured as a presupposition. Romoli (2013), for instance, argues that neg-raising should be treated as a kind of scalar implicature.

Let us now see what happens when we apply (86) to sentences of the form given in (92):

(92) Jack agrees with Sue on whether ϕ

Regarding the semantics of polar questions, recall that the strongy exhaustive answer to *whether* ϕ in world *w* is the proposition expressed by ϕ if ϕ is true in w, and its negation if ϕ is false in w. In the case of a whether-question $?\phi$, there is simply no difference between $[\![?\phi]\!](w)$, $exh_{?\phi}([\![?\phi]\!](w))$ and $Kart_{?\phi}([\![?\phi]\!](w))$,The set of potential complete answers to *whether* ϕ is therefore simply $\{[\![\phi]\!], [\![\neg\phi]\!]\}$ for all our notions of complete answers.

(86) predicts the following for (92):

(93) *Jack agrees with Sue on whether* ϕ presupposes that for some member p of $\{[\![\phi]\!], [\![\neg\phi]\!]\}$, Sue believes p and and asserts that for some member p meeting this condition, Jack believes p.

This is equivalent to:

(94) *Jack agrees with Sue on whether* ϕ presupposes that either Sue believes ϕ or Sue believes $\neg\phi$ and asserts that if Sue believes ϕ, then Jack shares this belief, while if Sue believes $\neg\phi$, then Jack shares that belief as well.

So (92) is predicted to presuppose that Sue has an opinion as to the truth-value of ϕ, and to assert that John shares this opinion.

When we turn to the negation of (92), 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* ϕ:

(95) *Jack does not agree with Sue on whether* ϕ presupposes that Sue is opinionated with respect to ϕ, and asserts that Jack does not share Sue's opinion about ϕ.

The sentence is thus predicted to be true if Sue, say, believes ϕ, and Jack has no opinion. However, as in the declarative case, we seem to observe something stronger, namely the negated sentence seems to imply that Jack actually believes that Sue's opinion is false. We assume here again that the relevant strengthening mechanism is not presuppositional (whatever its source is). If however we treated it as presuppositional from the start (i.e adopted a lexical entry for declarative agree that would include an opiniatedness presupposition regarding the subject of *agree with*), we would predict this inference as an automatic outcome of the presupposition.

So far, these predictions seem to be right. Importantly, it is quite clear that *X agrees/doesn't agree with Y on whether* ϕ does indeed presuppose that Y is opinionated with respect to ϕ, which shows that we want our lexical rule to be such that the presuppositions induced by *agree* when it takes a declarative complement are in some sense inherited when *agree* combines with an interrogative clause.

Let us now turn to a more complex example, in which the embedded question is a *wh*-question instead of a *whether*-question, as in (96). We will first determine what (86) predicts and then discuss the extent to which the predictions are correct, in light of previous discussions in the literature (Beck & Rullman 1999, Lahiri 2002, Sharvit 2002, George 2011) and a recent experimental paper by Chemla & George (Chemla & George, to appear).

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

(86) results into the following:

(97) **Presupposition**: (96) presupposes that for some potential strongly exhaustive answer S to 'Which students came to the party', the presupposition of *Jack agrees with Sue that S* is met.
Assertion: (96) asserts that for some potential strongly exhaustive 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:

(98) **Presupposition of (96)**: For some potential strongly exhaustive answer S1 to 'Which students came to the party', Sue believes S1.
Assertion of (96): For some potential strongly exhaustive answer S2 to 'which students came to the party' such that Sue believes S2, Jack believes S2.

It turns out that necessarily S1 and S2 are identical: indeed potential strongly exhaustive answers are mutually exclusive; thus since the presupposition entails that Sue believes a certain potential strongly exhaustive answer S1, it follows that if (96) is true, S1 is the only potential strongly exhaustive answer such that Sue believes it (assuming of course that Sue's beliefs are consistent). Hence (96) ends up asserting that for some potential strongly exhaustive answer S1, both Jack and Sue believe S1 is true.

Now, this prediction is different from what has been discussed in the previous literature. Lahiri (2002), taking inspiration from Hintikka (1976) and Berman (1991), takes (96) to simply mean that for every student *x* such that Sue has an opinion as to whether *x* came, Jack agrees with Sue, without any presupposition that Sue knows which students came in the strongly exhaustive sense. Specifically, Lahiri predicts (96) to be felicitous and true in a situation where Jack and Sue are uncertain about the same students and agree with each other about all the others. Beck & Rullman (1999) only discuss reciprocal *agree* (*Peter and Sue agree on which students came*), but they generally accept Lahiri's empirical claims. In a recent experimental survey based on a sentence-picture matching task (where agents' beliefs were represented as pictures), Chemla & George (2014) report results which are consistent with Lahiri's predictions.

What we would like to say (much in the spirit of Lahiri 2002 and following suggestions of George 2011), is that the reading we predict is the correct one, but that a mechanism of *presupposition-driven domain restriction* makes it possible to judge (96) felicitous and true in some situations where Jack and Sue are uncertain about some specific students.[24] The idea, which we will only present at an informal level, is the following. It is

[24] In the first version of our proposal, we did not take note of the fact that we predicted a reading that was stronger than what the previous literature had assumed. Chemla & George's recent survey prompted us to revisit our proposal. Many thanks to Ben George and Alexandre Cremers for extremely useful discussions.

quite generally possible to accommodate an underlying domain restriction so that a sentence's presupposition are met. For instance, a sentence such as *Every student in the school stopped smoking* can be felicitous even when it is known that some students have never smoked, because it can be reinterpreted as *Every student in the school who used to smoke stopped smoking*, i.e. the denotation of *student in the school* gets restricted to the set of students who meet the presuppositions of the predicate. Likewise, for the presupposition of (96) to be met, it is necessary that the underlying domain of individuals be such that Sue is opinionated about every student. So in case it is not already known whether Sue has an opinion about every student, we are led to posit a domain restriction such that the noun *students* ends up denoting the set of students about which Sue has an opinion – i.e. so that the presuppositions of (96) are met. Presupposition-driven domain restriction, we will assume, is always *minimal*, i.e. allows one to restrict the domain of individuals only as much as necessary for the relevant presuppositions to be met. Somewhat more explicitly, given a certain initial domain D of individuals, presupposition-driven domain restriction leads one to interpret the relevant sentence relative to the *maximal subset D' of D such that, relative to D', the presuppositions of the sentence are satisfied, where 'maximal' means 'that includes every other such set'* (there could fail to be a such maximal set in principle, and this will play a role in our discussion of reciprocal *agree*). This process leads to exactly the same predictions as Lahiri's. As Chemla & George note, no available theory matches exactly their data, but Lahiri's approach comes the closest.[25] Since our approach, supplemented with a mechanism of presupposition-driven domain restriction, makes the same prediction as Lahiri's (as presented by Chemla & George), it is at least as good as existing approaches. Let us note that the mechanism of presupposition-driven restriction is not expected to play a role in the case of *know*. This is so because in the case of *know+question*, as we have seen, the factive presupposition of *know* becomes trivial (i.e. is the tautology), so there cannot be any presupposition-related reason to restrict the domain of individuals.

If our account of *agree with* is on the right track, then if one can find a way to block presupposition-driven restriction, we should be able to detect the truth-conditions that we predict. One possible way to do this is to construct an example in which an explicit domain is made salient and important. Here is one attempt:

(99) *Mary is not sure about which of the ten students came. For some students, she has a definite opinion. For others, she has no idea. I can tell you this:*
Peter agrees with Mary about which of the ten students came.

We expect presupposition-driven accommodation to be quite hard in such a context (because the context makes the set of ten students highly salient and relevant), and as a result the sentence should not be fully felicitous given the immediate context. Specifically, (99) should be less felicitous than the following:

(100) *Mary feels certain about which of the ten students came. She has a definite opinion about all of them. I can tell you this*:
Peter agrees with Mary about which of the ten students came.

Our impression is that there is a contrast in the expected direction between these two cases, but we recognize that no firm conclusion can be reached without further investigation.

---

[25] Presupposition-driven domain restriction yields the consequence that *John agrees with Mary on which students came* and *Mary agrees with John on which students came* may have a different truth-value. The first sentence will mean *For every student about whom Mary has an opinion, John has the same opinion*, and this does not entail that for every student about whom John has an opinion, Mary has an opinion. Such an asymmetry, which is also predicted by Lahiri (2002), was not detected in Chemla & George's survey.

Let us now consider the 'reciprocal' use of *agree* as exemplified in (101) - (103)

(101) Jack and Sue agree that it is raining.
(102) Jack and Sue agree on whether it is raining.
(103) Jack and Sue agree on which students came.

What are the presuppositions triggered by reciprocal *agree* when it takes a declarative complement? To know this, we need to look at the negation of (101):

(104) Jack and Sue don't agree that it is raining.

It seems to us that this sentence suggests that one of Jack and Sue believes that it is raining and that the other does not share this believe. So we will assume that this sentence presupposes that at least one of Jack and Sue believes it's raining.[26] 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 (102), the presupposition we predict is that there is an answer *A* to *is it raining?* such that at least one of Jack and Sue believes A. And the assertion simply states Jack and Sue in fact both believe the answer A in question . Under negation (*Jack and Sue don't agree on whether it's raining*), the sentence is expected to mean that one of them has an opinion about whether it is raining that the other does not have (the other could be uncertain or actually disagrees). Such a prediction seems reasonable to us.

Things get more complex for (103). In this case, we predict the following:

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

So, in the absence of any presupposition-driven domain restriction, the sentence ends up entailing that both Jack and Sue believe a certain strongly exhaustive answer. However, allowing for presupposition-driven restriction, the domain will be restricted to the maximal domain such that either one of Jack and Sue is fully opinionated about this domain, if there is such a *maximal domain*. We can show that the sentence can only be true if Jack and Sue are opinionated *about exactly the same students* and are in full agreement about this set of students.

Let us see why. First suppose that there is no inclusion relation between the set of students about whom Jack is opinionated and the set of students about whom Sue is opinionated. Then there is simply no unique maximal set of students such that either Jack or Sue is opinionated about all the members of the set, and as a result presupposition-driven domain restriction cannot take place. Second, suppose that the set of students about whom Jack is opinionated is a proper subset of the set of students about whom Sue is opinionated.

---

[26] Assuming that *A and B agree that* ϕ should be analyzed as equivalent to *A and B agree with each other that* ϕ, one would need to derive this presupposition from the meaning of *agree with*, the reciprocal construction, and general principles of presupposition projection. This is far from trivial. If we analyze such sentences are equivalent to *A agrees with B that* ϕ *and B agrees with A that* ϕ, given that *X agrees with Y that* ϕ presupposes that *Y believes that* ϕ, we expect that *A and B agree that* ϕ will presuppose that both A and B believe ϕ, clearly a wrong result (because then such sentences could not be false, only true or undefined). This problem was noted by Lahiri. Just like us, Lahiri suggests a disjunctive presupposition 'either A or B believe that ϕ'.

Then the maximal set such that either Jack or Sue is opinionated about all members of the set is the set of students about whom Sue is opinionated. Presupposition-driven restriction thus leads us one to restrict the denotation of *students* to that set. But, relative to this set, Jack cannot believe a strongly exhaustive answer (since he is uncertain about at least one member of this set), and therefore the sentence is false. So the sentence can only be true if a) both Jack and Sue are opinionated about the same students, and b) have the same opinions about each of these students. This exactly corresponds to the truth-conditions predicted by Lahiri (2002) and discussed in the literature (Beck & Rullmann 1999, Sharvit 2002), and is also consistent with Chemla & George's experimental results.

V. Weak and Strong Exhaustivity, Intermediate Readings.

In this section, we propose a modification of our basic proposal, so as to be able to capture so-called weakly exhaustive readings. Let us clarify that in our final proposal, we will actually derive *intermediate readings* (to use Klinedinst & Rothschild's 2011 terminology) rather than weakly exhaustive readings in the standard sense. We are commited to the view that the examples that were used to argue for the existence of the weakly exhaustive readings are in fact instances of the intermediate readings. But we will first propose a rule that derives weakly exhaustive readings in the standard sense (a rule that we will amend later in order to generate intermediate readings instead). We will have nothing to say about the fact that certain predicates admit both the weakly/intermediate and the strongly exhaustive readings, while others have been said to be compatible with only one of them (George 2011 and Klinedinst & Rothschild 2011 claim that *know* is only compatible with the strongly exhaustive reading, but this is not what Chemla & Cremers found, and according to Guerzoni & Sharvit 2007 *know* licenses both the weakly and the strongly exhaustive reading, but *surprise* only supports the weakly exhaustive reading – but Klinedinst & Rothschild disagree). 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 (later *intermediate)* 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, strongly exhaustive and intermediate readings.

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

Consider the following sentence:

(106) 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 that nobody else came (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, Guerzoni & Sharvit 2007 for arguments for, and Cremers & Chemla, to appear, for experimental evidence that something close to the weakly exhaustive reading exists). One argument for the existence of the weak reading is based on the fact that the following is not deviant (Guerzoni & Sharvit 2007):

(107) 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 (107) 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 (107) 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 be the case, (107) starts sounding like a contradiction. That is, does (108) sound contradictory?

(108) 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 (108) 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.[27]

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 that were initially pointed out by G&S and were then discussed again 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

[27] See George (2011) for an extensive discussion. George concludes that neither the weakly exhaustive reading nor the intermediate reading exists for *know*. Results from recent experimental surveys do not support this conclusion (Rothschild & Klinedinst 2011 report the results of such a survey for the verb *predict*, but not *know*, and Cremers & Chemla, to appear, provide more systematic evidence).

know which of the guests attended the party.[28] 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, G&S themselves briefly considered a slighly stronger characterization that does not run into this problem, which was subsequently discussed in Spector (2005, 2006) (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. A recent paper by Cremers & Chemla (Cremers & Chemla, to appear) provides strong experimental evidence that such a reading (which they call, after Klinedinst & Rothschild 2011, the 'intermediate reading') exists, controlling for the potential role played by domain restriction. This is the reading we call the *intermediate reading*.

Under this intermediate reading of embedded questions, (107) would therefore be satisfiable, even when the domain of quantification is constant and known to John. 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). We return to this in section V.4.

To conclude, our approach so far is not able to derive the weakly exhaustive reading that other approaches derive. This in itself 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. However there seems to exist a slightly different reading, which we call the intermediate rading, 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 return to this issue in section VII. In the next sections we will show that we need to be able to derive weakly or intermediate exhaustive readings for at least some embedding verbs, and we will propose a modification of our proposal that does capture such readings.[29]

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

---

[28] 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.

[29] As Groenendijk & Stokhof (1982) noticed, on the weakly exhaustive reading *know* does not satisfy positive introspection when it embeds an interrogative clause. That is, John can know who came without knowing that he knows who came. This is the case if for every *x* who came John knows that *x* came and for every *y* who didn't come John has no idea whether *y* came.

so-called emotive predicates (*amaze*) favor a weakly exhaustive reading. Consider for instance the following sentence.

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

Imagine 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, (109) 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.[30] That is, (109) 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 (109).

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 base the theory on strongly exhaustive compete answers (i.e. *GS(Q)*), we get the strongly exhaustive reading. But if we use a weaker notion, as shown in section III.1, we don't get the weakly exhaustive reading, but a much weaker reading (the tautology if we refer to basic complete answers).

However, in the subsequent sections, we will make 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.

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

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 some weaker

[30] 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). George (2011), on the other hand, argues that the weakly exhaustive reading does not exist with *surprise*, and that the cases that are used to claim otherwise are in fact instances of the mention-some reading. See our discussion in footnote 34.

notion. We will discuss in section VIII and in the appendix which of the two 'weaker' notions we introduced (basic complete answers or Karttunen-complete anwers) should be used, but this is not relevant at this point. These two notions differ from each other only in the special case where the question-predicate has an empty extension (say in case nobody came when the question is 'Who came?'). So we will first ignore these slight differences and just use the basic notion of complete answer, until section V.6.. Let us illustrate our idea 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.[31]

(110) It surprises Mary that S
→ Presupposition: S and Mary believes S (= Mary knows S)

What our proposal will derive for *It surprises Mary Q*, where *Q* is a constituent question, is the following (informally):

(111) **Presupposition:** There is a potential *basic* complete answer *p* such that the presuppositions associated with *It surprised Mary* $exh_Q(p)$ are satisfied.
Assertion: There is a potential *basic* complete answer *p* such that a) the presuppositions associated with *It surprised Mary* $exh_Q(p)$ are satisfied, and b) *p* is surprising to Mary.

Let us see what this amounts to in the case of (109). The predicted presupposition is that there is a potential strongly exhaustive 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 strongly exhaustive answer that is true in the actual world is the actual strongly exhaustive answer, this amounts to saying that the actual strongly exhaustive answer is true (a tautology) and that Mary knows that it's true. Because strongly exhaustive 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 strongly exhaustive answer. Hence the predicted presupposition amounts to *Mary knows which of the guests showed up*, *in the strongly exhaustive sense*.

Let us now turn to the assertive part. By 'assertive' part, we simply mean here a characterization of the truth value of the sentence *in every world where its presuppositions are satisfied*, which is what the right hand-side of our lexical rules does. We do not assume here a bi-dimensional theory of presupposition in which one could refer to the non-presuppostional content of a sentence as a separate dimension of meaning.[32] The assertive part

[31] The assumption that *be surprising to X that S* presupposes knowledge of *S* by *X* is not uncontroversial. Egré (2008) argues that so-called emotive factive predicates, like *regret*, are not in fact veridical, but involve only the weaker presupposition that *X believes p*. Here we choose to abide by the more standard assumption that *surprise* is indeed factive, and comes with a knowledge presupposition. Nothing essential hinges on this.

[32] In a trivalent approach to presuppositions, there is no need to define a non-presuppositional 'assertive component' for a presuppositional expression. For instance, *John knows that S* denotes a partial function from worlds to propositions, such that the function is not defined for a world in which *S* is false. There is no non-arbitrary way (and no need) to extract from this partial function a complete function which would correspond to a non-presuppositional, bivalent 'assertive' component. This is so because there are many different ways of 'completing' a partial function, and so no non-arbitrary way to define such a non-presuppositional, bivalent proposition that would express the 'assertive part' of *John knows that S*. Of course, in our metalanguage, when we state the truth-conditional import of an expression, we need to make use (in the right-hand side of our rules) of a metalanguage statement which we can informally call the 'assertive part'. This part tells us how to assign truth and falsity to a sentence *whose presuppositions are satisfied*, and there are many equivalent ways to formulate it which can look superficially very different. For instance, the three following (simplistic) lexical entries for *know* are fully equivalent:

in this sense says that there is a potential basic complete answer *p* such that $exh_Q(p)$ is true and Mary knows that $exh_Q(p)$ is true, and furthemore *p* is surprising to Mary. Now, again, the only *p* that can meet this condition is the *actual* basic complete answer. So the assertion can be reformulated as follows: *Mary knows the truth of the actual strongly exhaustive complete answer, and she is surprised by the actual basic 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 (109) 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:

(112) Context: Mary has 10 students, and they all took a certain exam. She definitely did not expect students A, B and C to pass, but 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, they know 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.[33] That is, even though Mary is surprised by what is, *in fact*, the basic 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.

Note, however, that the mechanism of presupposition-driven domain restriction discussed in relation with *agree* in section IV.2 should allow one to interpret the sentence in such a way that its presuppositions are satisfied, by restricting the domain of individuals to the set of individuals such that Mary knows who among them came. It is plausible that the context we set up in the above example interferes with this mechanism and makes it quite costly.

(i) $[\![know_{decl}]\!]^w = \lambda p.\lambda x$: p(w) = 1. x believes p in w
(ii) $[\![know_{decl}]\!]^w = \lambda p.\lambda x$: p(w) = 1. p(w) = 0 or x believes p in w
(iii) $[\![know_{decl}]\!]^w = \lambda p.\lambda x$: p(w) = 1. p(w) = 1 and x believes p in w

[33] 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.

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, at this point, the following (again, recall that we do not discuss yet which specific notion of weak complete answer should be used, so this proposal is going to be revised later).

(113) 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_{weak}, has the following lexical entry:
$[\![P_{weak}]\!]^w = \lambda Q.\lambda x : \exists p \in pot(Q) \text{ s.t. } [\![P_{decl}]\!]^w(exh_Q(p)) \text{ is defined.}$
$\exists p \in pot(Q)\ ([\![P_{decl}]\!]^w(exh_Q(p)) \text{ is defined } \ \&\ [\![P_{decl}]\!]^w(p)(x) = 1)$

Let us see what this yields for (109) ('It surprised Mary which guests showed up'). First, the definedness condition yields the presupposition that Mary knows the truth of the actual strongly exhaustive answer. This is so because we assume that *It surprises X that A* is not defined if A is not true. So a potential basic answer *p* can satisfy the existential statement of the definedness condition only if $exh_Q(p)$ is true, and there is, in a given world *w*, a *unique* such basic complete answer, namely the actual basic answer. Because *It surprises X that A* also presupposes that *X* knows *A*, we get the presupposition that Mary knows $exh_Q(p)$, where *p* is the actual basic complete answer, i.e. that Mary knows the truth of the actual strongly exhaustive answer. According to the right part of the rule (the 'assertive' part), (109) is true if and only if Mary knows which guests showed up in the strongly exhaustive sense, and Mary is surprised by the actual *basic* complete answer, i.e. by the 'positive' part of the strongly exhaustive answer. This is the desired weakly exhaustive reading.

Let us comment on our strategy for deriving weakly exhaustive readings. 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)[34] is to make sure that the potential basic complete answer which makes true the existential statement "There is a potential weakly exhaustive answer A such that x is related to A via the attitude relation P" corresponds to a

[34] One could wonder whether a very weak semantics in terms of basic complete answers, of the type we rejected in section III.1, could 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* should be felicitous and true. But according to our informants, for the sentence to be felicitous, Ann has to know which guests attended the party. One possible fix would be to claim that at the presuppositional level the strong notion of complete answer is used, with the results that *It surprised Ann which of the guests came* would presuppose that Ann knows which of the guests came (in the strong sense) and that for some guests who came it surprised her that they came. According to our intuitions, however, 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 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 S* entails *it surprised Mary that S & T*, which seems natural at first sight but might not be correct after closer scrutiny).

strongly exhaustive answer $exh_Q(A)$ which is such that the presuppositions of 'x V $exh_Q(A)$' are true.

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. On the basis of (113), we get the following, where *w* is the world of evaluation.

(114) Mary knows which guests showed up

a. **Presupposition**: there is a potential basic complete answer A to 'which guests showed up?' such that $[[know]]^w(exh_Q(A))(Mary)$ is defined, i.e. such that $exh_Q(A)$ is true.

b. **Assertion**: there is a potential basic complete answer A to 'which guests showed up?" such that a) $exh_Q(A)$ is true and b) $[[know]]^w(A)(Mary) = 1$

As before, the presupposition is trivial: it just says that there is a basic potential complete that its associated strongly exhaustive answer is true. This is always the case, the actual basic complete answer in the world of evaluation is such that its associated strongly exhaustive answer is true (and is in fact the unique true strongly exhaustive answer). Let us now turn to the assertion. It says that there is a potential basic answer A such that a) its strongly exhaustive answer is true, and b) Mary knows the truth of A. As we have just seen, only the *actual* basic complete answer in *w* can meet the a)-condition. So the assertion ends up equivalent to the claim that Mary knows the truth of the *actual* basic complete answer. Since the presupposition, as we have just seen, is trivial, this is nearly identical to the reading predicted by Karttunen. The only difference comes from the special case where the basic complete answer is the tautology (i.e. when no guests showed up). In this case, the actual basic complete answer is the tautology, and so we predict that (114) is true if no guest showed up and Mary knows that the tautology is true, which might well be itself a tautological statement. This is clearly a problem, which we now address

V.5. Using Karttunen-answers rather than basic complete answers

We present here a specific proposal regarding the treatment of tautological basic complete answers, but there are several reasonable approaches, all of which are compatible with our general goals. Because the treatment of this specific case (i.e. the case where the actual basic answer is the tautology) is not the main focus of our paper, we defer the discussion of some of these variants to an appendix.

As it stands, our proposal makes a clearly undesirable prediction: that in a world where nobody came and Peter has no knowledge whatsoever, *Peter knows who came* is true under the weakly exhaustive reading. This is so, because in this situation Peter happens to know the actual basic complete answer, namely the tautology.

Karttunen offered a solution by defining a complete answer as being equivalent to a basic complete answer except in the special case where the basic complete answer is a tautology – in this case, Karttunen wanted the complete answer in his sense to be equivalent to the strongly exhaustive complete answer. So a natural way out of this problem is to follow Karttunen's footsteps and replace *p* with $Kart_Q(p)$ in the assertive part of our rule in (113). So the rule for weakly exhaustive readings becomes the following:

(115) 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_{weak}, has the following lexical entry:
$[\![P_{weak}]\!]^w = \lambda Q.\lambda x : \exists p \in pot(Q) \text{ s.t. } [\![P_{decl}]\!]^w(exh_Q(p)) \text{ is defined.}$
$\exists p \in pot(Q)\ ([\![P_{decl}]\!]^w(exh_Q(p)) \text{ is defined } \ \&\ [\![P_{decl}]\!]^w(Kart_Q(p))(x) = 1)$

For *John knows which students showed up*, the prediction we now make is that, in case no student showed up, the sentence is true if and only if John knows that no student showed up.

For *It surprised Mary which students showed up*, the prediction now is that the sentence presupposes that Mary knows which students showed up in the strongly exhaustive sense, and that either no student shows up and this fact surprised Mary, or Mary is surprised by the 'positive' part of the strongly exhaustive answer to *which students showed up*. It is not clear to us whether this is a good prediction. Our impression is that *It surprised Mary which students showed up* can hardly be used to describe a situation where Mary was surprised that no students showed up, but we are not quite sure about this, and don't know of any compelling evidence one way or the other. So we will adopt (115) as the basis for our account of weakly exhaustive answers.[35]

V.6. Interim conclusion

Our proposals for weakly exhaustive readings and strongly exhaustive readings are parallel. They are identical for the presuppositional part, and differ only regarding one parameter, namely the choice of the relevant notion of complete answer. Our proposal can therefore be stated in a compact way as follows:

(116) **Meaning postulate for both the strongly and the weakly exhaustive reading (provisional)**
$[\![P_{strong/weak}]\!]^w =$
$\lambda Q.\lambda x{:}\exists p \in pot(Q)\ ([\![P_{decl}]\!]^w(exh_Q(p))(x) \text{ is defined}).$
$\exists p \in pot(Q)\ ([\![P_{decl}]\!]^w(exh_Q(p)) \text{ is defined } \&\ [\![P_{decl}]\!]^w(\mathbf{param}_{\mathbf{Q}}(p))(x)=1),$
where **param** is a variable that takes as its value either *exh* or *Kart*.

How does this proposal fare? We managed to solve the problem that we noted in section III, which prevented us from using a weaker notion of complete answer than that of the strongly exhaustive 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). By making use of *both* strongly exhaustive answers and K-complete answers 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. In section VII, we will tentatively replace our rule for deriving weakly exhaustive readings with a rule that will derive *intermediate readings*, along the lines of Klinedinst & Rothschild (2011). But we will first address another issue that we have not touched upon yet, which will already motivate a slight change of perspective.

---

[35] We may want such a sentence to actually count as a presuppositional failure in a situation where no guests showed up. See the appendix, where this possibility is explored.

VI. Taking into account presuppositional monotonicity.

So far we were dealing with cases where the presuppositions associated with a predicate were, in an intuitive sense, *positive*. More formally, these were cases where the presuppositional content of the predicate can be described as *monotone increasing*. Let us define the function *Presup* as a function which, when applied to a sentence, returns the presupposition of the sentence. The point is that (informally) *Presup(X knows that p) = p* and *Presup(It surprised X that p) =(p and X believes p)*. In both cases, whenever we have *p => p',* we also have *Presup(....p...) => Presup(...p'...)*. This is what allowed us to characterize a weakly exhaustive reading that was not *too* weak (cf. our discussion in section III).

But there exist predicates whose assertive part is *positive*, but whose presuppositions are negative or at least non-monotonic. It turns out that such predicates create a problem for our characterization of the *strongly* exhaustive readings, and probably also for weakly exhaustive readings [Danny Fox (p.c.) and Jeroen Groenendijk (p.c.) both drew our attention to such predicates. See in particular Fox (2013) for a sketch of a proposal that is related to the one that we present here, which is based on a suggestion of Jeroen Groenendijk.]

A case in point is the verb *discover*. A sentence such as *Mary discovered that A* presupposes not only that *A* is true, but also that at some point before the time of the discovery, Mary didn't know *A*. So the lexical entry for *discover* can be informally described as follows (we need to add a *time* coordinate as an evaluation parameter, and of course all compositional rules would have to be modified accordingly):

(117) $[\![\text{discover}]\!]^{w,t} = \lambda p.\lambda x$: p(w) = 1 and there is a t'< t such that x did not know p in t' in w. there is a t'< t such that x did not know p in t' in w and x believes p in w.

We will focus now on this 'epistemic' presupposition when *discover* embeds a question, as in (118):

(118) Mary discovered yesterday which guests showed up.

Consider the following scenario: in fact the guests who came are guests A and B. At some point before yesterday Mary knew that A and B showed up, and had no idea about the other guests. Yesterday Mary learnt that *only* A and B showed up. On the basis of our rule for strongly exhaustive readings as defined in (116) (by picking *param* = *exh*), (118) comes out true in this scenario. This is so because the predicted presupposition is that Mary didn't know before yesterday the truth of the actual strongly exhaustive answer (which is the case in this scenario), and the assertion is that Mary yesterday came to believe the actual strongly exhaustive answer (which is also the case in this scenario). However, it seems to us (following D. Fox and J. Groenendijk, to whom we owe this observation) that in such a scenario the sentence is not true, and might well be, in fact, a presupposition failure. In a nutshell, (118) appears to presuppose that Mary didn't know before yesterday the *positive* answer to the question, i.e. the actual basic complete answer. In any case, the way we defined the strongly exhaustive reading runs into a problem with this case.

Note that in this scenario the sentence comes out false (not undefined) if we use the weakly exhaustive reading as defined in (116) (by picking *param* = *Kart*), equivalent to (115). This is so because given that Mary, in this scenario, didn't know before yesterday the actual complete answer, and so the presupposition is satisfied. But since she knew before yesterday the truth of the actual basic complete answer (which is in this case the actual K-complete answer) the assertive part itself is not satisfied: in this scenario, $[\![\text{discover}_{\text{decl}}]\!]^{w}(\text{Kart}_{Q}(p))(x)$, where p is the actual basic complete answer, is undefined, and so the condition in the assertive part '$[\![\text{discover}_{decl}]\!]^w(\text{Kart}_Q(p))(x)=1$' cannot be met (Mary cannot have discovered p if she already knew p). So on the basis of (115), the sentence is predicted to presuppose that Mary didn't know before yesterday the actual strongly exhaustive answer, and to *assert* that she also didn't know the actual K-complete answer and that she then came to believe it. This is a better result, but still the sentence is not predicted to *presuppose* that Mary didn't know the 'positive' part of the actual strongly exhaustive answer. And yet (118) does seem to presuppose that Mary didn't know this 'positive part'.

What seems to be the case is that with *discover* the definedness condition should be stated in terms of basic complete answers, rather than in terms of strongly exhaustive answers, both for the weakly and the strongly exhaustive readings. We would then have the following.[36]

(119) **Strongly exhaustive reading for *discover***
$[\![\text{discover}_{strong}]\!]^w = \lambda Q.\lambda x{:}\exists p \in \text{pot}(Q)([\![\text{discover}_{decl}]\!]^w(p)(x) \text{ is defined}).$
$\exists p \in \text{pot}(Q)\ [\![\text{discover}_{decl}]\!]^w(\text{exh}_Q(p))(x) = 1.$

(120) **Weakly exhaustive reading for *discover***
$[\![\text{discover}_{weak}]\!]^w = \lambda Q.\lambda x{:}\exists p \in \text{pot}(Q)([\![\text{discover}_{decl}]\!]^w(p)(x) \text{ is defined}).$
$\exists p{\in}\text{pot}(Q)([\![\text{discover}_{decl}]\!]^w(\text{exh}_Q(p))(x) \text{ is defined} \ \&\ [\![\text{discover}_{decl}]\!]^w(\text{Kart}_Q(p))(x){=}1)$

On the basis of (119), (118) is predicted to presuppose that there is a potential basic complete answer that is true and whose truth Mary didn't know, i.e. that for at least one guest who actually showed up,[37] Mary didn't know that this guest showed up, and asserts that there is a potential strongly exhaustive answer that Mary discovered, i.e. a) is true, b) was not known by Mary before yesterday, and c) is now known by Mary.[38] This seems to us to be the most salient reading for (118). An interesting consequence, which we view as desirable, is that the sentence is predicted to be a presupposition failure in case no guests showed up. This is so because if no guests showed up, the basic complete answer is the tautology, and as a result the presupposition that Mary didn't know the basic complete answer is contradictory – i.e., if no guests showed up, the presupposition cannot be met. In other words, (118) is predicted to presuppose that at least one guest showed up. Note that we get this result from the fact that in the presuppositional part we still use the basic notion of complete answer instead of Karttunen-answers (cf. footnotes 36 and 37).[39]

On the basis of (120), (118) is again predicted to presuppose that there is a potential basic complete answer that is true and whose truth Mary didn't know. And it is predicted to assert that Mary now knows the truth of the actual K-complete answer. This weakly exhaustive reading, we believe is available as well (putting aside the *no-false belief* problem discussed in section V, to which we will provide a solution in section VII). Note that in the assertive part, we crucially need the first conjunct which refers to strongly exhaustive answers, so as to ensure that the K-complete answer that Mary now knows is not only true, but is the actual K-complete answer (recall that a K-complete answer can be true without

[36] Note that in the presuppositional part we still use the basic notion of complete answer rather than Karttunen-answers. See the appendix, where we discuss the possibility of not making use at all of basic complete answers.
[37] If no guest showed up, the basic complete answer is the tautology and the presupposition that Mary did not know the basic complete answer is necessarily false.
[38] Recall that we are using a trivalent theory of presuppositions where if a sentence is true, then its presuppositions are true as well.
[39] There may be a tension between this prediction and the fact that for *It suprised Mary who came*, we do not predict presuppositional failure in case nobody came. See footnote 35 and our discussion in the appendix.

being the actual K-complete answer). The reader is invited to check that similar lexical entries would be needed for other predicates whose presuppositions are not monotonic, such as, e.g., *realize*.

What is going on here? It looks as if the presuppositions of *V+interrogative* can be stated in terms of strongly exhaustive answers if *V*'s presuppositions are monotone increasing in the sense defined above, but not otherwise. As pointed out to us by Danny Fox and Jeroen Groenendijk, we can make sense of this by formulating the definedness condition as a conjunction, so that each conjunct refers to a different notion of complete answer. When the presupposition induced by the predicate is monotone increasing, the conjunct that includes a reference to strongly exhaustive answers happens to entail the other one, so that this other conjunct plays no role. This results in the following modification of our lexical rules (again, following ideas put forward by Danny Fox and Jeroen Groenendijk).

(121) **Meaning postulate for the strongly exhaustive reading (revised)**
$[\![P_{strong}]\!]^w =$
$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined).
$\exists p \in pot(Q)\ [\![P_{decl}]\!]^w(exh_Q(p))(x) = 1)$

(122) **Meaning postulate for the weakly exhaustive reading (revised)**
$[\![P_{weak}]\!]^w =$
$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined). $\exists p \in pot(Q)\ ([\![P_{decl}]\!]^w(exh_Q(p))$ is defined & $[\![P_{decl}]\!]^w(Kart_Q(p))(x) = 1)$

To maximize the parallelism between the two rules and provide a uniform format where the presuppositional part is repeated in the assertive part (so that it can properly *constrain* the identity of the basic answer that will satisfy the existential statement of the assertive part), we can propose the following general rules.

(123) **Meaning postulate for the strongly exhaustive reading (revised)**
$[\![P_{strong}]\!]^w =$
$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined).
$\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)=1)$

(124) **Meaning postulate for the weakly exhaustive reading (revised)**
$[\![P_{weak}]\!]^w =$
$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined).
$\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))$ is defined & $[\![P_{decl}]\!]^w(Kart_Q(p))(x)=1)$

As far as we can see, for all the cases we have discussed, repeating the presuppositional part in its entirety in the assertive part makes no difference (but it could in principle).[40] This seems to us to be the most motivated option: conceptually, the idea we have been capitalizing on is that the presuppositional part serves as a constraint on how the existential condition of the assertive part can be instantiated. This allows us to present these two rules in a compact form, as follows.

---

[40] For instance, repeating '$[\![P_{decl}]\!]^w(p)$ is defined' in the assertive part of the rule for strongly exhaustive readings generally has an effect for a predicate whose presuppositional content is not monotone increasing with respect to its complement. In particular, if P is presuppositionally monotone-decreasing, '$[\![P_{decl}]\!]^w(p)$ is defined' entails '$[\![P_{decl}]\!]^w(exh_Q(p)$ is defined', and its addition is thus non-vacuous. We haven't seen such a case yet. Even though *discover* is presuppositionally non-monotonic, the condition '$[\![discover_{decl}]\!]^w(exh_Q(p)$ is defined' is sufficient to force p to be the unique actual basic complete, because *discover* is factive.

(125) **Meaning postulate for both the strongly and weak exhaustive readings (revised)**

$[\![P_{strong/weak}]\!]^w =$

$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x) \text{ is defined } \& [\![P_{decl}]\!]^w(exh_Q(p))(x) \text{ is defined}).$

$\exists p \in pot(Q)([\![P_{decl}]\!]^w(p) \text{ is defined } \& [\![P_{decl}]\!]^w(exh_Q(p)) \text{ is defined } \& [\![P_{decl}]\!]^w(\mathbf{param_Q}(p))(x)=1),$

where **param** is a variable that takes as its value either *exh* or *Kart*.

VII. Deriving the 'No false-belief' constraint (intermediate readings in Klinedinst & Rothschild's 2011 sense).

We pointed out in section V.1 that 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 (in the case of *know* and some other verbs), briefly discussed by Groenendijk and Stokhof, and then in Preuss (2001) and Spector (2005) among others. 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. As already emphasized, we are not aware of any fully satisfying account that would solve this problem without stipulations specifically tailored to solve it. We can however offer a solution which is partly inspired by Klinedinst & Rothschild (2011).

One ingredient of our solution is somewhat theoretically costly. We will need to adopt a 'decompositional' approach to presuppositional attitude predicates, according to which the lexical entry for such verbs specifies not only a presuppositional content (say, through a definedness condition), but also an 'assertive' content that can be neatly separated from the presuppositional content and is not itself presuppositional (such a move is already briefly suggested by Klinedinst & Rothschild (2011), though they ultimately refrained from it and assumed that *know* only licenses strongly exhaustive readings).[41] For instance, *know that S* would be decomposed into a presuppositional part (*S* is true) and a non-factive attitude whose meaning might be approximated by *believe*.[42] From now on, we will assume that it is possible to extract from a presuppositional attitude predicate P a non-presuppositional, assertive part, which we will simply note $[\![)\, P_{decl} (]\!]$. Lexical elements thus now receive two values, the standard value represented by the double-bracket notation, and the 'non-presuppositional' value represented by using the modified double-bracket notation ($[\![)$ $(]\!]$). As an illustration, the lexical entry for *know* might look like the following:

(126) a. $[\![)\, know_{decl} (]\!]^w = \lambda p.\lambda x.\ x \text{ believes } p \text{ in } w.$

b. $[\![know_{decl}]\!]^w = \lambda p.\lambda x{:}p(w)=1.\ [\![)\, know_{decl} (]\!]^w(p)(x)=1$

---

[41] As discussed in footnote 32, in a trivalent approach to presuppositions there is no intrinsic need to be able to extract a non-presuppositional 'assertive' part from a presuppositional expression, even though we informally refered to the right-hand side of our lexical rules as to an 'assertive component'. Our lexical entries, which were of the form $\lambda x.\lambda y. \ldots .\lambda z{:}\varphi(x,y,\ldots,z).\psi(x,y,\ldots,z)$, are just notational devices to describe partial functions, with the lefthand side specifying when the function is defined, and the righthand side specifying which value the function returns when it is defined.

[42] We totally ignore here the fact that *true belief* is not sufficent to characterize *knowledge*. A more accurate characterization of the non-factive attitude associated with *know* may be *believe on the basis of compelling subjective evidence*. This however still runs into Gettier's problem (Gettier 1963). Note that the idea that it is possible to extract from *know* a corresponding non-factive attitude has been argued to be intrinsically misguided on the basis of variations on Gettier's problem – see in particular Williamson 2002.

A possible approach to the no-false belief problem then consists in replacing our rule for the weakly exhaustive reading with a rule that captures *intermediate readings* (in the sense of Klinedinst & Rothschild 2011), i.e. a rule that incorporates the no-false belief constraint:

(127) **Meaning postulate for the intermediate reading (provisional)**
$[\![P_{intermediate}]\!]^w$ =
$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined).
$\exists p \in pot(Q)$ ($[\![P_{decl}]\!]^w(p)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))$ is defined & ($[\![P_{decl}]\!]^w(Kart_Q(p))(x)$ = 1 & **there is no p'$\in$ pot(Q) such that both p' asymmetrically entails p and** $[\![)\mathbf{P}_{\mathbf{decl}}(]\!]^w$**(p')(x)=1)**

Let us see what happens in the case of *know*. At the level of the presupposition, nothing changes. But for *John knows Q* to be true, it must now be the case not only that John knows the truth of the actual K-complete answer to Q – call it $Kart_Q(A)$ -, but also that there is no other potential complete answer B to Q such that B asymmetricaly entails A and John *believes* that B is true. If *Q* is *Who came?*, this ensures that there is no person *x* such that John erroneously believes that *x* came. It is crucial here that the additional clause is stated in terms of $[\![)P_{decl}(]\!]$ rather than in terms of $[\![P_{decl}]\!]$. In the case of *know*, if the additional clause refered to $[\![P_{decl}]\!]$, it would be trivially satisfied. This is so because any potential basic complete answer B that asymetrically entails the actual basic complete answer A is necessarily false, and since *know* is factive (hence veridical), it is thus trivially the case that for every such B it is not the case that John knows B, even if John *believes* B – but it is not trivial to say that for every such B it is not the case that John believes B.

Yet this modification runs into problems for predicates whose monotonicity properties are different from that of *know*. Consider for instance *forget*.[43] *X forgot S* presupposes that *S* is true, that *X* used to know *S*, and asserts that *X* does not currently believe *S*. The corresponding non-factive attitude that would capture the assertive part of *forget* is then something like *not-believe*, i.e. is motonone-decreasing with respect to is complement. Then for *John forgot which guests showed up*, (127) results in a near-contradiction. The sentence would be predicted to be true if, on the one hand, John does not currently believe the truth of the actual K-complete answer $Kart_Q(A)$, but, on the other hand, believes every potential basic complete answer B that asymmetrically entails A. Now note that if *X* is in the relation *not-believe* to a proposition p, X is necessarily in the relation *not-believe* to any stronger proposition. So, if $Kart_Q(A)$ = A (as in the case if at least one guest showed up) and X forgot $Kart_Q(A)$, X is necessarily in the relation *not-believe* to any proposition B that entails A, and therefore the condition that John believes every potential complete answer that is stronger than A cannot be met. So the sentence is predicted to be false whenever $Kart_Q(A)$ = A, hence is predicted to *entail* that no guests showed up (and, specifically, to assert that John forgot that no guests showed up and now believes that every guest showed up!) – clearly a wrong result.[44]

Klinedinst & Rothschild (2011) propose an approach that can probably be adapted to provide a general rule that works for all these cases. Here we offer a proposal that is different

---

[43] Thanks to Alexandre Cremers for drawing our attention to this case.

[44] One possibility that suggests itself would be to replace, in (127), the clause 'such that p' asymmetrically entails p' with 'such that $\lambda v.[\![)P_{decl}(]\!]^v(p')(x) = 1 \subset \lambda v.[\![)P_{decl}(]\!]^v(p)(x)$', where $\subset$ represents *asymmetric* entailment. This, however, doesn't solve the problem for *John forgot which guests showed up*. Suppose, for instance, that in fact the people who came are John and Mary. Let us schematize 'John and Mary came' with *j & m.* Now note that *X forgot j* entails *X forgot j & m*, and so does *X forgot m* (in this case, whether or not we consider only the assertive part or include the presupposition does not matter). Then with the modifed clause, (127) would result in the following truth conditions (when the people who came are *John and Mary*) for *Peter forgot who came*: Peter forgot *j & m* but neither forgot *j* nor forgot *m*, which is clearly contradictory.

and is easier to present (though maybe one that is also more *ad hoc*).[45] Our goal is only to show that a solution is *possible*. Our solution is that in the additional clause that takes care of the no-false belief constraint, we now quantify over alternative potential basic complete answers *p'* such that <u>at the same time</u> *p'* a-symmetrically entails *p* and *x V p'* assymmetrically entails *x V p*. The point is that when *V* is not monotone increasing, this conjunction is generally a contradiction, i.e. no potential basic complete answer can make both conjuncts true. That is, our proposal is the following:

(128) **Meaning postulate for the intermediate reading (revised)**
$[\![P_{intermediate}]\!]^w =$
$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined).
$\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))$ is defined & $[\![P_{decl}]\!]^w(Kart_Q(p))(x)=1$
& there is no p'∈ pot(Q) such that
$\mathbf{p' \subset p}$
& $\boldsymbol{\lambda v.[\![P_{decl}]\!]^v(p')(x)=1 \subset \lambda v.[\![P_{decl}]\!]^v(p)(x)=1}$
& $\boldsymbol{[\![P_{decl}]\!]^w(p')(x)=1}$)

For *Mary knows which guests showed up*, the clauses $\mathbf{p' \subset p}$ and $\boldsymbol{\lambda v.[\![P_{decl}]\!]^v(p')(x)=1 \subset \lambda v.[\![P_{decl}]\!]^v(p)(x)=1}$ are equivalent (because $\boldsymbol{[\![know]\!]}$ – i.e. *believe* – is monotone increasing) But for *Mary forgot which guests showed up*, the point is there is no pair (p,p') that can simultaneously satisfy both conditions and so the additional clause that takes care of the no-false belief constraint (in boldface) is simply vacuously true for *forget*. In the case of *It surprised Mary who came*, we get a similar result. Assume that the assertive part of *surprise* is equivalent to *was not expected*. There is generally no pair of propositions (p, p') such that at the same time p' a-symmetrically entails p and *It was not expected by Mary that p'* entails *it was not expected by Mary that p*. Indeed, if p' is strictly stronger than p, it is certainly *possible* for Mary not to expect p' and at the same time to expect p, and so *it is not expected by Mary that p'* does not entail *it is not expected by Mary that p*. This is because *surprise* is not monotone increasing with respect to its complement. In other words, for *surprise*, the intermediate reading and the standard weakly exhaustive readings end up being equivalent. [46]

Let us briefly see what is predicted for a few other question-embedding verbs. Consider *discover*. Assuming (which is certainly simplistic) that the assertive part of *discover* is simply *believe*, we get the same result as for *know*. Namely, for *John discovered who came* to be true on the intermediate reading, John must not falsely believe that a certain person came who in fact didn't come. The predictions for *agree* are more complex. *Recall* that *Mary agrees with Peter that S* presupposes that Peter believes S, and asserts that Mary believes S.

[45] Klinedinst & Rothschild derive the no-false belief constraint by applying an exhaustivity operator to the V+Q constituent. Very informally, the idea is this: 'X V Q' is true if, a) X is in the relation V to the actual basic complete answer to Q – call it A and b) the *exhaustification* of the proposition 'X V A' relative to alternative propositions of the form X V B, where B can be any potential complete answer to Q, is true as well. Because exhaustification is defined in such a way that only *innocently excludable alternatives* in Fox's (2007) sense are negated, it can never result in a contradiction. This approach can presumably be adapted to our own framework.

[46] We do not take a stand as to whether *surprise* is non-monotonic with respect to its complement or monotone decreasing. Let us note that if *surprise* is treated as monotone-decreasing with respect to its complement (for instance by stating that *p is surprising to X* if and only *X* assigned *p* a probability lower than some threshold before *X* learnt *p*), Klinedinst & Rotschild's proposal runs into a problem. It predicts that *it surprised Mary who came* is true if a) Mary was surprised that the people who came came and b) for any subplurality of the people who came, Mary was not surprised that this plurality came. For instance, if the people who came are *a*, *b*, and *c*, Mary has to be surprised that *a* and *b* and *c* came, but not that *a* and *b* came, not that *a* and *c* came and not that *b* and *c* came. This prediction does not seem correct to us. Klinedinst & Rothschild thus need to assume that *surprise* is in fact non-monotonic (which we think might well be true). Thanks to Nathan Klinedinst for relevant discussions.

As a consequence, the assertive part of *agree with Peter* is simply *believe*. (128) then results in the following, informally.

(129) Interpretation of *Mary agrees with Peter on who came* under (128)
Presupposition: there is a potential basic complete answer A such that Peter believes exh(A)
Assertion: Mary believes the potential K-complete answer A that makes the presupposition true, and Mary does not believe any stronger potential basic complete answer.

Now, assuming that Peter is coherent and that the presuppositions of the sentence are met, there is one and only one potential basic complete answer A such that Peter believes exh(A). Then Mary has to believe the 'positive part' of the strongly exhaustive answer that Peter believes, and should not have any other 'positive' belief. Allowing for presupposition-driven domain restriction, where the domain of individuals is restricted to the largest domain relative to which the presuppositions of the sentence are met, the prediction is that the sentence is true if, for every individual such that Peter believes that this individual came, Mary agrees with Peter, and for every individual such that Peter believes that this individual did not come, Mary does not believe that this individual came. Whether such a reading exists is not clear to us (it is not detected as such in Chemla & George's survey).

In the preceding section, we proposed very similar rules for the weakly and strongly exhaustive readings, which we viewed as a desirable feature. Can we add the additional clause that takes care of the *no-false belief constraint* to the lexical rule that captures strongly exhaustive reading, without any bad consequence? The answer is yes, because for all the cases we have considered this additional clause is vacuous, something that we leave it to the reader to verify.

So we are in a position to state our final proposal in terms of two completely parallel rules which differ from each other only in terms of the choice of a specific parameter (which notion of complete answer is used in the assertive part). These two rules can thus be presented as one rule that includes a variable:

(130) **Meaning postulate for both the strongly exhaustive and the intermediate readings (final)**

$[\![P_{strong/intermediate}]\!]^w =$

$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined).

$\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x)$ is defined & $[\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined & $[\![P_{decl}]\!]^w(\mathbf{param}_Q(p))(x)=1$

& there is no $p' \in pot(Q)$ such that

$p' \subset p$

& $\lambda v.[\![P_{decl}]\!]^v(p')(x)=1 \subset \lambda v.[\![P_{decl}]\!]^v(p)(x)=1$

& $[\![P_{decl}]\!]^w(p')(x)=1)$

where **param** is a variable that takes as its value either *exh* or *Kart*.

In many cases, however, these postulates are reducible to more simple statements. First, if we pick **param** = exh, then the part that takes care of the 'no-false' belief constraint is, for all the predicates we have checked, vacuous. Second, if the predicate P is presuppositionally monotone increasing, we can eliminate all the occurrences of "$[\![P_{decl}]\!]^w(p)$ is defined", because this is then always entailed by "$[\![P_{decl}]\!]^w(exh(p))$ is defined".

Note also that is we pick **param** = exh, since we have "$[\![P_{decl}]\!]^w(exh_Q(p))(x)=1$" entails "$[\![P_{decl}]\!]^w(exh_Q(p))$ is defined", we can also eliminate the part "$[\![P_{decl}]\!]^w(exh_Q(p))$ is defined".

With predicates such as *forget*, which are Strawson-decreasing (i.e. decreasing in their assertive part: 'forget' = 'fail to believe now') but presuppostionally increasing (e.g., 'forget' presupposes 'used to know'), (130) reduces to the following (because the part that takes care of the 'no false-belief' constraint is now vacuous):

(131) **If P is Straswon-decreasing:**
$[\![P_{strong/intermediate}]\!]^w$ =
$\lambda Q.\lambda x{:}\exists p\ [\![P_{decl}]\!]^w(exh_Q(p))(x)$ is defined).
$\exists p{\in}pot(Q)([\![P_{decl}]\!]^w(exh_Q(p))$ is defined & $[\![P_{decl}]\!]^w(\mathbf{param}_Q(p))(x)=1$
where **param** is a variable that takes as its value either *exh* or *Kart*.

In these cases, the intermediate reading as we defined it is equivalent to the standard weakly exhaustive reading.

VIII. 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 *intermediate readings* of embedded questions (a modified version of the so-called *weakly exhaustive* reading), 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 intermediate reading. 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. We would like to make clear what we think we have achieved. We have provided *general* meaning postulates such that the meaning of the interrogative-taking variant of a responsive predicate is predictable (modulo the *weak/strong* distinction) from the meaning of the declarative-taking variant. However, given that we have said nothing about what determines which of the intermediate and strongly exhaustive readings are available, our proposal does not entirely eliminate lexical stipulations. Lexical stipulations might be needed to state that some predicates but not others license, say, the weakly exhaustive reading. Nevertheless, our proposal goes beyond existing proposals, in narrowing down to two the number of possible meanings for the *V+interrogative* construction, once the meaning of the *V+declarative* construction is known.

Our theory is however not *explanatory* in the following sense: it does not derive the meaning postulates we posited from first principles. As pointed out by a reviewer, the ingredients of our proposal would allow us to define different meaning postulates on top of those we posited. For instance, we characterize the intermediate reading of the interrogative-taking variants in such a way that, at the presuppositional level, the strong notion of complete answer is used, while the weak notion is used at the assertive level. One could easily posit another, 'symmetric', meaning postulate whereby the interrogative-taking variant would involve the weak notion of complete answer at the presuppositional level, and the strong notion at the assertive level. Our theory does not *explain* why we need the meaning postulates that we posit rather than other ones. Our contribution is thus to have shown what an explanatory theory (which is yet to come) would have to derive.

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 (see Egré 2008 for a tentative proposal).[47] Similarly, we would like to understand why certain verbs, e.g. *surprise*, can take a wh-question as a complement but not a polar question (see Guerzoni 2007 for a proposal). These are topics for further research.

Appendix – More on tautological basic answers

There are various ways in which our proposal might be modified regarding its treatment of the special case where the actual basic complete answer is the tautology. In this appendix, we discuss three possible modifications of our proposal.

A.1 Eliminating any reference to basic complete answers.

Note that our proposal makes use of three distinct notions of answers: basic complete answers, Karttunen-complete answers, and strongly exhaustive answers, where Karttunen-complete answers and basic complete answers only differ regarding the special case where the basic complete answer is the tautology. Now, we can wonder what would follow if we eliminated completely any reference to basic complete answers, i.e. if we used Karttunen answers only, even in the presuppositional part of our meaning postulates. In such a case we could simply define *Q(w)* differently, so that it would directly denote what is currently denoted by $Kart_Q(Q(w))$. In other words, *pot(Q)* would now be identifed to the set of potential

[47] 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.

Karttunen-complete answers, rather than that of basic complete answers. But instead of introducing new definitions, let us rather keep our definitions, in terms of which we can state the resulting proposal as follows:

(132) **Meaning postulate for both the strongly exhaustive and the intermediate readings (modified version #1)**

$[\![P_{strong/intermediate}]\!]^{w} =$

$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^{w}(Kart_{Q}(p))(x)$ is defined & $[\![P_{decl}]\!]^{w}(exh_{Q}(p))(x)$ is defined).

$\exists p \in pot(Q)([\![P_{decl}]\!]^{w}(Kart_{Q}(p))(x)$ is defined & $[\![P_{decl}]\!]^{w}(exh_{Q}(p))(x)$ is defined & $[\![P_{decl}]\!]^{w}(\mathbf{param}_{Q}(p))(x)=1$

& there is no $p' \in pot(Q)$ such that

$p' \subset p$

& $\lambda v.[\![P_{decl}]\!]^{v}(p')(x)=1 \subset \lambda v.[\![P_{decl}]\!]^{v}(p)(x)=1$

& $[\![P_{decl}]\!]^{w}(p')(x)=1)$

where **param** is a variable that takes as its value either *exh* or *Kart*.

Now, compared to our ‘official’ proposal, this does not change anything for predicates which are ‘presuppositionally’ monotone-increasing, such as *know*, *forget*, and *surprise* (*X knows that S* presupposes *S*, *X forgets that S* and *It surprised X that S* presuppose *S and X knows/used to know S*). In particular, we still predict that *It surprised Mary who came*, on the intermediate reading, is neither a presuppositional failure nor trivally false in case nobody came. That is, such a sentence will entail that if nobody came, Mary knows it and is surprised by that fact. The proposal however makes a difference for predicates which are presuppositionally non-monotone-increasing, such as *discover* (*X discovers that S* presupposes *S and X did not believe S*). Whereas on the ‘official’ proposal, as we discussed, *John discovered who came* is a presupposition failure if in fact nobody came, on this new proposal the sentence can be true if nobody came, provided John didn’t know at some point that nobody came and came to believe that nobody came. We tend to think that our ‘official’ proposal makes a better prediction, but, as pointed out by a reviewer, it might be problematic not to make parallel predictions for *surprise* and for *discover*. That is, one might think that if *John discovered who came* is a presupposition failure when nobody came, so should be *It surprised John who came*.

A.2 A presuppostional variant for the weakly exhaustive reading.

Another possibility would be to introduce, only in the rule for the intermediate reading, a presupposition ensuring that the basic complete answer to which the agent is related via the relevant attitude is not the tautology. On such a view, the rule for the strongly exhaustive reading would remain the same as in our official proposal, but we could not state the rule for the intermediate reading in terms of just a change of ‘parameter’ in the assertive part of the rule. We would also distinguish both readings in terms of their presuppositional behavior. The rule for the intermediate reading would then be stated as follows:

(133) **Meaning postulate for the intermediate readings (modified version #2)**

$[\![P_{intermediate}]\!]^{w} =$

$\lambda Q.\lambda x{:}\exists p \in pot(Q)(p$ is not the tautology and $[\![P_{decl}]\!]^{w}(p)(x)$ is defined & $[\![P_{decl}]\!]^{w}(exh_{Q}(p))(x)$ is defined).

$\exists p \in pot(Q)(p$ is not the tautology and $[\![P_{decl}]\!]^{w}(p)(x)$ is defined & $[\![P_{decl}]\!]^{w}(exh_{Q}(p))(x)$ is defined & $[\![P_{decl}]\!]^{w}(p)(x)=1$ & there is no $p' \in pot(Q)$ such that

$p' \subset p$

& $\lambda v.[\![P_{decl}]\!]^{v}(p')(x)=1 \subset \lambda v.[\![P_{decl}]\!]^{v}(p)(x)=1$

& $[\![P_{decl}]\!]^{w}(p')(x)=1)$

If we adopt this proposal, both sentences *John knows who came* and *It suprised John who came*, on the intermediate reading, would presuppose that someone came. The fact that *John knows who came* is not judged to be a presuppositional failure, but is rather judged true, if nobody came and John knows that fact, would then be attributed to the availability of the strongly exhaustive reading (which would not trigger such a presuppostion). It might also be expected that *It surprised Mary who came* would tend to be perceived as infelicitous if it is known that nobody came, given that the intermediate reading might be the prefered reading with *surprise* (in fact, Guerzoni & Sharvit 2007 argue that it is the *only* reading). Nothing would change for *John discovered who came*. It would still be predicted to be a presuppositional failure if nobody came, both on the intermediate and on the strongly exhaustive reading (cf. our discussion in section VI).

A3 – A more radical move: eliminating any reference to Karttunen-complete answers

Recall the motivation for using Karttunen-complete answers rather than basic complete answers for the intermediate reading: we do not want to predict that *John knows who came* is true when nobody came and John has no belief whatsoever. But we may try to explain this intuition not in terms of the *semantics* of embedded questions, but in pragmatic terms. A potential motivation for using basic complete answers rather than Karttunen-complete answers has to to with verbs such as *tell* and *predict* on their non-veridical use, and with non-presuppositional question-embedding predicates such as *be certain about*. In a nutshell, we will show such a move allows us to predict an interesting generalization: on their non-factive, non-veridical uses, verbs like *tell* and *predict* only license a strongly exhaustive reading. And likewise a non-presuppositional predicate like *be certain about* is only compatible with the strongly exhaustive reading.

Let us first illustrate the generalization. For the sentence *Mary is certain about which students came* to be true, Mary must be certain about a certain strongly exhaustive answer. If she is only certain, say, that the students Sue and Al came and has no idea for others, the sentence is false. As for *predict*, a sentence such as (134) below entails that Mary made a *complete* prediction, in the sense that her prediction was interpreted as a claim about what the actual strongly exhaustive reading is [Of course Mary does not need to have explicitly predicted who would come and who would not come. She may have said 'Mary and Peter will come', which, as an answer to 'Who will come', is easily intepreted as implying that nobody besides Mary and Peter will come].

(134) Peter predicted which of the four students would attend the party, but he proved wrong.

To see this, consider (134) uttered in the context described in (135) below:

(135) a. <u>Scenario:</u> 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 made no prediction about the others – he actually said that he had no idea about Fred and Mary. In fact it turned out that neither John nor Sue attended the party.
b. <u>Sentence:</u> Peter predicted which of the four students would attend the party, but he proved wrong.

According to our informants, (135)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.

(136) 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 (135) and (136), in that (136), unlike (135), can be judged true. This shows that the non-veridical use of *predict*, which may be less salient than its veridical use, is in any case only compatible with the strongly exhaustive reading.

Now, it turns out that this state of affairs is easy to predict if we adopt the proposal in (137) instead of our official proposal, whereby the intermediate reading is defined with no reference to Karttunen-complete answers (note that this proposal is essentially what you get if you add to (113) the clauses that take care of the 'no-false belief' constraint).

(137) **Meaning postulate for the intermediate readings (modified version #3)**

$[\![P_{intermediate}]\!]^w =$

$\lambda Q.\lambda x{:}\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x) \text{ is defined } \& [\![P_{decl}]\!]^w(exh_Q(p))(x) \text{ is defined}).$

$\exists p \in pot(Q)([\![P_{decl}]\!]^w(p)(x) \text{ is defined } \& [\![P_{decl}]\!]^w(exh_Q(p))(x) \text{ is defined } \& [\![P_{decl}]\!]^w(p)(x)=1$

& there is no $p' \in pot(Q)$ such that

$p' \subset p$

$\& \lambda v.[\![P_{decl}]\!]^v(p')(x)=1 \subset \lambda v.[\![P_{decl}]\!]^v(p)(x)=1$

$\& [\![P_{decl}]\!]^w(p')(x)=1)$

On this proposal, we get the following prediction:

(138) Mary predicted which guests attended the party.

Presupposition: there is a potential basic complete answer p to 'which guests attended the party?' such that $[\![predict]\!]^w(p)(Mary)$ and $[\![predict]\!]^w(exh_Q(p))(Mary)$ are defined.

→ This is always true (because non-factive *predict* has no presupposition)

Assertion: there is a potential basic complete answer p to 'which guests attended the party?" such that Mary predicted p and Mary did not predict any stronger basic complete answer.

The presupposition is vacuous. As to the assertion, it is true as soon as Mary predicted that *some* potential basic complete answer is true. Take indeed the *strongest* potential basic complete answer that Mary predicted. Such a proposition satisfies the existential statement of the assertive part of the rule. So even if the strongest potential basic answer that Mary predicted is the tautology, the sentence comes out true. In fact, the sentence ends up equivalent to *Mary predicted the tautology*. Assuming that *predict* is monotone-increasing with respect to its complement, as soon as Mary predicted something, she predicted the tautology. So the sentence is true when Mary predicted the tautology and entails that Mary predicted the tautology, hence it simply expresses the proposition that Mary predicted the tautology. Likewise, if we use non-veridical, non-factive *tell*, *Mary told Peter who came* comes out true if and only if Mary told Peter the tautology. This, we surmise, is sufficient to explain why the weakly exhaustive reading is not detected for these non-presuppositional uses of *tell* and *predict*. On some quite standard assumptions *X told Y the tautology* and *X predicted the tautology* are themseleves tautologies. Even if we do not subscribe to such a view, the interpretation that results from our rule is such that the content of the embedded clause plays no role at all in the truth conditions, which may be sufficient to rule out this reading on pragmatic grounds. Finally, with *be certain about*, on the weakly exhaustive reading, *John is certain about who came* would come out equivalent to *John is certain that the tautology is true*, which is plausibly itself tautological.

It is important to note here that with *know*, or any other presuppositional predicate, the rule for intermediate readings as stated in (137) does not give rise to such trivial readings, thanks to the presuppositional part of the rule, which becomes non-vacuous for these verbs and constrains the identity of the basic complete answer that can satisfy the existential statement. *Mary knows which students came* is predicted to be equivalent (on the intermediate reading) to 'Either no student came and and there is no student x such that Mary believes that x came, or Mary knows, for every student who came, that these students came, and does not falsely believe, of a student who didn't come, that he came'. But if we adopt (137), we still need to explain why, on the intermediate reading, *Mary knows which students came* is not judged true when Mary does not have any idea about who came and in fact no students came (this is the problem of tautological basic answers).

We would like to suggest that this judgment might itself be explained in pragmatic terms. Note that when it is common knowledge that no student came, *Mary knows which students came* becomes equivalent to the claim there is no student such that Mary believes that this student came (but she may have no belief at all about students). So in such a context, the sentence does not assign *any* knowledge to Mary. Though we cannot pin down the precise underlying pragmatic principle that would make such a use of a *know* deviant, it seems to us that a pragmatic explanation for this specific case is plausible – for instance, because the use of *know* in *Mary knows Q* normally licenses the inference that Mary knows something about Q. This line of explanation actually predicts that we should be able to set up a context where the sentence is in fact judged true when Mary has no knowledge whatsoever. What we need is a context where the speaker does not have the belief that no student came, but does not exclude that possibility either. On top of that, the relevant scenario must be such that the attitude bearer (the subject of *know*) is presented as having no false positive beliefs about the question but also as having no true negative beliefs (so that we are sure that we probe the intermediate reading). Constructing such a context and then the appropriate scenario proves pretty hard, which is in itself suggestive, as it may explain the intuition that if Mary has no knowledge whatsover *Mary knows Q* is false (in fact only quite implausible contexts and scenarios seem to meet all these conditions). Nevertheless, here is an attempt to set up such a context (thanks to Danny Fox for suggesting this specific scenario):

(139) Every Friday at 5PM, the radio gives the names of the people who won the lottery. John's father, who is old and has memory problems, listens to the radio all the time. If nobody wins, John's father doesn't notice that no list is read, so he does not know that nobody wins. If a list is read, he doesn't know that it's the complete list. But he does remember the list of the names for some time after 5PM.
On a certain day, John tells his wife at 5.05pm: "I'll call my father. **He knows who won the lottery**"

If in fact on that day nobody won the lottery, was John's sentence false? Our impression (and that of our informants) is that even upon learning that there was no winner and so John's father knew nothing at all, one does not conclude that John was wrong. We conclude, very tentatively, that it *might* be possible to defend the proposal (137).

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.

Ciardelli, I. (2009), Inquistive Semantics and intermediate logics. Master Thesis, U. of Amsterdam.

Ciardelli & Roelofsen, F. (2009), Generalized inquisitive logic: Completeness via intuitionistic Kripke models. *Proceedings of Theoretical Aspects of Rationality and Knowledge.*

Ciardelli & Roelofsen, F. (2011), Inquisitive Logic. *Journal of Philosophical Logic*, 40(1): 55-94.

Ciardelli, I. and Roelofsen, F. “Inquisitive Dynamic Epistemic Logic”. To appear in *Synthese*.

Ciardelli, I and Groenendijk, J. and Roelofsen, F. (2013). “Inquisive Semantics : a new notion of meaning”. *Language and Linguistics Compass* 7(9), 459-476, 2013.

Chemla E. & George B. (2014). “Can we agree about *agree*”? Ms.

Cremers, A. and Chemla, A. (2014), A psycholinguistic study of the different readings for embedded questions. Ms. Available at http://semanticsarchive.net/Archive/DU3YWU2M/cremerschemla-semarch.html

Dayal, V. (1996). *Locality in WH quantification: Questions and relative clauses in Hindi*. Dordrecht: Kluwer Academic Publishers.

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.

Fox, D. (2013), *Classnotes on mention-some readings*., Ms., MIT, Available at http://web.mit.edu/linguistics/people/faculty/fox/class1-3.pdf

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. (2009), Inquisitive semantics: Two possibilities for disjunction. In P. Bosch, D. Gabelaia & J. Lang (eds), *Seventh Tblisi Symposium on Language, Logic and Computation*, Springer-Verlag.

Groenendijk, J. & Roelofsen, F. (2009), Inquisitive Semantics and Pragmatics. Presented at the Workshop on Language, Coommunication and Rational Agency, Stanford 2009. URL www.illc.uva.nl/inquisitivesemantics.

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.

Hamblin, C. L. (1973). Questions in Montague english. *Foundations of language*, 41-53.

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.

Hintikka, J. (1976). *The Semantics of Questions and the Questions of Semantics*, Cornell University Press, Ithaca, NY.

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.

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

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.

Mascarenhas, S. (2009). *Inquisitive Semantics and Logic*. MSc in Logic thesis. Amsterdam: Institute for Logic, Language, and Computation.

Preuss, S. (2001). *Issues in the Semantics of Questions with Quantifiers*, PhD dissertation, Rutgers.

Roelofsen, F. (2013), Algebraic foundations for the semantic treatment of inquisitive content. *Syntehese,* 190(1) : 79-102, doi : 10.1007/s11229-013-0282-4.

Roelofsen, F., Theiler, N & Aloni, M. (2014), Embedded interrogatives : the role of false answers. Presented at the 7th Questions in Discourse Workshop, Göttingen, Sept. 2014.

Romoli, J. (2013). A scalar implicature-based approach to neg-raising. *Linguistics and philosophy*, *36*(4), 291-353.

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

Schlenker P. (2007), "Transparency: An Incremental Theory of Presupposition Projection", in U. Sauerland & P. Stateva (eds.), *Presupposition and Impicature in Compositional Semantics*, Palgrave: NY, p. 214-242.

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*, PhD dissertation, University of 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

Theiler, N. (2014), A multitude of answers: embedded questions in typed inquisitive semantics. MSc thesis, University of Amsterdam, supervised by Maria Aloni and Floris Roelofsen.

Tsohatzidis, S. L. (1993), "Speaking of Truth-Telling: the View from *Wh*-complements", *Journal of Pragmatics* 19, 271–279.

Tsohatzidis, S. L. (1997). More telling examples: A reply to Holton. *Journal of Pragmatics* 28, 629-536

Williamson, T. (2002), *Knowledge and its limits*. OUP