

# Quantified exhaustiveness in embedded questions

Ang Li    *Rutgers University*

May, 2018

## 1 Introduction

The topic of this paper is about deriving the correct truth conditions of sentences like (1):

- (1) The principal mostly catches who cheat on the final exam.
- (2) The principal usually catches who cheat on the final exam.

Since Berman (1991), It has become a well-accepted observation that the sentence in (1) has a reading different from (2)<sup>1</sup>. In particular, while (2) talks about the frequency of catching *all* cheaters in final exams, (1) seems to have a reading where the adverb quantifies over the wh-expression(s) in the embedded clause. Suppose the principal catches 75% of the cheaters in every final exam, only (1) is clearly true. This has been discussed under the term “Quantificational Variability Effects” (QVE). The problem posed by this interpretation is how the adverb finds its domain of quantification, and what exactly the domain is.

A lot of proposals have been made (most prominently Berman (1991), Groenendijk & Stokhof (1992), Lahiri (2002), Beck & Sharvit (2002)), and this paper will add yet another one. I differ from all these previous attempts in suggesting that QVE is simply verb modification. The impression that the quantificational force interacts with the embedded question is made by having the adverb attaching to a “resolution” layer inside the verb. In other words, QVE is about quantifying over exhaustiveness, but interestingly the data seems to suggest that the exhaustiveness are inherently derived from the embedding verb, not the question itself.

---

<sup>1</sup>Actually Berman’s original work uses *usually* to demonstrate QVE, but Lahiri (2002) shows only amount quantifiers are reliable triggers for this reading, frequency adverbs are not.

The structure of the paper is as follows. Section 2 introduces the problems of previous accounts, focusing on their inability to derive the sensitivity to false answers. Section 3 elaborates on this point to show that the theoretical problems lies in taking the quantificational domain of the adverb to be the semantic objects of the embedded question. Section 4 lays out my proposal, based on the observation that some verbs are resolute. Section 5 shows how the verb classification naturally accounts for the restriction on verb types of QVE, and finally gives some remarks on why the QVE data might be more interesting than the scope of this paper.

## 2 False-answer sensitivity in QVE

In his dissertation, Berman paraphrased sentences like (1) as roughly the following: “for most of the cheaters, the principal catches them on the final exam”. This paraphrase has been *mostly* well-accepted in the literature<sup>2</sup> and has been the basis of many formal accounts. Our key observation here, however, is that this paraphrase is incomplete, if not wrong. In particular, this paraphrase isn’t sensitive to false answers the subject might have. Suppose the principal catches the majority of *true* cheaters but also identifies a lot of innocent students as cheaters – in this case the sentence seems to be infelicitous. This doesn’t seem to be a pragmatic ignorance, either. Suppose we are told the principal catches most of the cheaters but also wrongly identifies one or two innocent students as cheaters sometimes; now the sentence “the principal mostly catches who cheat on the final exam” sounds like a legit judgement. In comparison to the situation before, it seems we *are* aware of the false answers, and the *portion* of the false answers matters in granting the *mostly* description.

The problem is even more outstanding if we consider verbs like *discuss*, *decide*, and *depend on*. When the embedding verb is one of those, the most salient QVE reading we get is clearly not restricted to true answers:

- (3) a. John and Mary mostly discussed who will be admitted.  
= For most candidates, John and Mary have discussed on whether they will be admitted.  
(Berman’s paraphrase: for most people who actually will be admitted, John and Mary discussed that they will be admitted.)
- b. John mostly decided who will be the attend the meeting.  
= For most people who might attend the meeting, John has decided

---

<sup>2</sup>But see Groenedjik & Stokhof (1992), Ginzburg (1995a), and Beck & Sharvit (2002).

whether they will or not.

(Berman's paraphrase: for most of people who actually will attend the meeting, John has decided that they will attend the meeting.)

In these cases, the paraphrase in Berman's way sounds odd, and are certainly not the most salient reading. What's going on?

I think the way out is actually very clear in its direction: we need to derive a reading that can look into the whole situation. The lack of false-answer sensitivity and the weird paraphrases is a result of we restricting the quantificational domain to true answers - or whatever subset of the positive possible answers<sup>3</sup>. The real challenge for semanticists is how this can be done technically and compositionally.

### 3 Quantification over the embedded question

Of course, Berman's paraphrase is well-accepted for a reason: that seems to be the most straightforward result, if you take one crucial assumption. That assumption is the adverb like *mostly* quantifies over the semantic objects of the embedded question; in other words, it's *mostly* accepted<sup>4</sup> that the QVE reading has a logical form like the following (I use  $p$  as a type-neutral variable here)<sup>5</sup>:

$$(4) \quad \llbracket \text{John mostly knows who walks} \rrbracket \\ = \mathbf{most}(\{p \mid p \in \llbracket \text{who walks} \rrbracket\}) (\{p \mid \mathbf{j}\text{-knows}(p)\})$$

What you take  $p$  to be therefore becomes the crux of the question, and has led to the discussion of what a quantificational domain the embedded question can produce. It can't be the set of possible answers (i.e. the Hamblin set), otherwise the "John mostly knows who walks" would have to mean "John *believes* most of the possible answers (are true)." If you care about the factive presupposition of *know*, this sentence may not even mean anything when the portion of true answers is small relative to the whole range of possible answers. That's why we are restricting ourselves to the Karttunen set (i.e. the subset of true answers; or other subsets, for verbs like *agree on*). While this provides an appropriate quantificational domain and a (sometimes) acceptable reading, we've seen that it's incomplete. The price we pay is losing the global view of the situation; what we need is to recover the negative extension of the question answers.

<sup>3</sup>Noticing the truth requirement does varies depending on the embedding verb, Lahiri (2002) does suggest the quantificational domain of the adverb can be some other subset of the Hamblin set.

<sup>4</sup>Again there are proposal who don't adopt this view, most prominently see Ginzburg (1995b).

<sup>5</sup>This is actually the logical form proposed by Lahiri (2002). Berman's solution is slightly different, where he takes the adverb to directly quantify over the indefinite variable, but the variable further gets restricted by the answer of the question.

Recovering the negative extension is usually attributed to the strongly exhaustive reading of an embedded question, in which a complete answer to the question “who walks?” not only specifies the true answers “*a* and *b* walk”, but also negates the false answer “*c* walks”. The problem of incorporating the strongly exhaustive reading into QVE is that, unfortunately, it does not provide a well-defined quantificational domain. For all of the theoretical options of getting the strongly exhaustive reading – world partition (Groenendijk & Stokhof (1989)), strong answerhood operator (Heim (1994)), or covert exhaustivity (Klinedinst & Rothschild (2011)) – the result amounts to a single proposition (suppose *w* is the real world, and only *a* and *b* walk in *w*):

- (5)  $\llbracket \text{who came} \rrbracket^w$  (strongly exhaustive reading)  
 $\equiv \lambda v \forall p \in \{\text{came } a, \text{came } b, \text{came } c\} : pw = pv$   
 $\equiv \lambda v (\text{came}_v a \wedge \text{came}_v b \wedge \neg \text{came}_v c)$

This clearly can’t be the domain of the adverb, or *mostly* will pick up *most* of the possible worlds which are identical to the real one in view of the question. That’s not the reading we are after, and probably not a possible reading at all.

What we really need, to get the “global view” we are after, is to recover not just the negative extension as a strengthened inference (which is what currently the strongly exhaustive reading offers), but each of the *negative* possible answers. Eventually we want a set of positive true answers and a set of negative false answers, and their intersection provides a maximally informative answer:

- (6)  $\llbracket \text{who came} \rrbracket^w$  (the ideal quantificational domain)  
 $\rightsquigarrow \{\text{came } a, \text{came } b, \neg \text{came } c\}$

The problem, again, is how this could be done in a principled way, because this certainly can’t be compositionally derived either above the Hamblin set or the Karttunen set. Beck & Sharvit’s (2002) attempt is, as far as I see, the most successful one in this regard. They define an operation that yields part polar questions out of a constituent question. In other words, the idea is this operation partitions over a question into a set of subquestions that checks each individual answers one by one:

- (7)  $\text{Part}(\llbracket \text{who came} \rrbracket^w)$   
 $\equiv \{\text{Did } a \text{ come? Did } b \text{ come? Did } c \text{ come?}\}$

Then the QVE reading amounts to say that “out of all these subquestions, John knows/ decides on most of them”. This gives the global view as we want, and false-answer sensitivity naturally follows. Nevertheless, this view makes it mysterious

why QVE is restricted to only a subset of question embedding verbs. Verbs like *wonder* and *ask* don't have the QVE reading at least in its standard form:

- (8) John mostly wonders which kids are drunk.  
≠ For most  $x$ ,  $x$  is a kid: John wonders if  $x$  is drunk.

The contrast between *wonder* and *know* can be nicely captured in Lahiri's (2002) account where he takes the latter to be only proposition-embedding and the QVE interpretation is derived through raising the embedded question - which is triggered by a type mismatch. This argument doesn't seem to work if the adverb quantifies over a set of subquestions, because now *know* has to also embed (sub)questions as well.

To summarize the situation we are in, the challenge in deriving the correct truth conditions comes back to the challenge of getting a well-defined quantificational domain out of the embedded question. One sensible way to recover the negative false answers is to *keep* the domain as a domain of *questions* in a sense, but in doing so the restriction on verbs becomes mysterious. Although I have to say the final point is not so much an argument *against* the view of subquestions - maybe what we need is just an additional theory on verb selection - it's always worth exploring whether we can have a theory of QVE that makes sense of all the relevant pieces.

#### 4 A new proposal

In this dilemma, it's worth thinking whether the crucial assumption - that QVE arises when the adverb quantifies over the semantic objects inside the embedded question - is warranted. Especially, consider the fact that so far the most successful attempt involves keeping the embedded question as *questions* instead of answers. I think it's clear that what we need is to *partitioning* over a total answer; what's not so clear but possibly a promising direction is my following proposal: what *mostly* does is simply changing the answerhood condition from universal quantification to the degree of *most*.

The official proposal comes in three pieces: i) QVE reflects (a special kind of) verb modification; ii) QVE arises when the embedding verb is *resolutive*; iii) the answerhood condition can be re-engineered as a resolution to a definite answer, and the nature of QVE is simply that the resolution is gradable.

The idea that QVE is simply verb modification has already been brought up by Ginzburg (1995b). However, instead of further considering the difference between the QVE reading and normal adverb modification, Ginzburg seems to prefer to

mash them all into some kind of degree modification and attributes the difference into the complement clause – QVE arises when the complement clauses are *questions* which are coerced into *facts*. This then raises a number of objections in other people’s reviewing. Here I agree with Ginzburg’s basic intuition, but I also need to stress that QVE is a *different* kind of modification; it modifies a *resolution* process that only some verbs have.

The second crucial idea is then to point out that *some* verbs are resolute. I’ll give a more explicit definition later, but it’s also quite intuitive to think that *know* and *discover* are different from *wonder* and *ask* in this regard. While *wonder* and *ask* care only about a certain mental state at one point, *know* and *discover* contains something more: they involve justification of one’s belief or wondering. To represent this in a formal talk, we may think that resolute verbs such as *know* and *discover* has an additional layer of meaning, as sketched in the decomposition in Figure 1. The decomposition can roughly be understood as first forming a

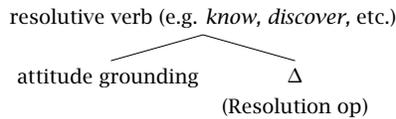


Figure 1: decomposing resolute verbs

belief state in the step of attitude grounding, and then justifying all your relevant beliefs are true in the step of resolution. At the end of the day, the meaning we have for *know* is roughly “having certain beliefs that are guaranteed to be true”, which seems to be on the right track.

Why this is interesting is because we can now see that the second layer that I call “resolution” really resembles what people have been discussing under the term of “answerhood conditions” or “answerhood operator” a lot. The final piece is exactly this: we can re-engineer the answerhood conditions in the form of a resolution operator, and let the verb combine with the Hamblin set directly.

Before entering into the formal details, I need to mention that this view has a conceptual advantage: it naturally captures the fact that *which* subset of the possible answers we want to extract out of the Hamblin set in the embedding environment is decided by the lexical semantics of the verb. It is defaulted to “truth” in a lot of cases, but certainly “truth” is not a generic requirement on embedded questions when you consider verbs that are neither factive or veridical:

- (9) a. John is certain who was at the party.

b. John and Mary agree who was at the party.

Neither (9a) nor (9b) entails that if Bill and Sue were at the party John should be certain about their presence or that he and Mary should agree about their presence. Rather, these sentences say if “Bill and Sue were at the party” is a belief that John has, he is certain about it; or in (9b) if “Bill and Sue were at the party” is a belief of either John and Mary, it is in the common ground between the two of them as well. In the standard theory of question-embedding, this is done by giving some flexibility to the answerhood conditions: restricting the Hamblin set by a variable  $C$  or having an evaluation world which is not necessarily the true world. However, the value of  $C$  or the evaluation world has to be said to depend on the *context*, whereas it’s clear that we are witnessing something different than normal context-dependency.  $C$  or the evaluation world should pick up something determined by the lexical semantics of the verb, but it’s hard to see how this could be done when the answerhood condition is separated from the verb but is attached to the embedded question. Now if all we have been calling as answerhood conditions is really a resolution layer encoded by the verb, all the dependencies are straightforward.

Now let’s see how we can derive these compositionally. I take the attitude grounding layer to simply denote an event of having an attitude, for *know* it is an event of believing. I follow Hacquard (2010) to assume that attitude event has propositional content, and the intersection of these propositions is the set of worlds compatible with his attitude (in this case what the subject believes). This can be retrieved out of a given event using her **CON** function: given an attitude event  $e$ , “ $\cap\text{CON}(e)$ ” is the set of all possible worlds compatible with the attitude. What the resolution operator does is gluing the event with a question together, by stating the existence of such an event, and that all of the possible world we retrieve using **CON** agree with some evaluation world on the judgement of the possible answers *to some degree*. In the case of *know* (and other factive verbs), the evaluation world is set to the real world. And in the unmarked case, the *degree* of resolution is defaulted to *all*. Therefore the underlying “resolution” we have can be formally represented like the following formula in (10), and you can imagine in the unmarked cases the quantifier place holder  $\kappa$  is saturated by a silent universal quantifier:

$$(10) \quad \Delta: \lambda\kappa_{a-t} \lambda f_{v-t} \lambda Q_{s-t}. \exists e f(e) \wedge \forall v \in (\cap\text{CON}(e)) \kappa p \in Q \ p w = p v$$

The full derivation of the sentence “John knows who walks” is given in figure 2 below.

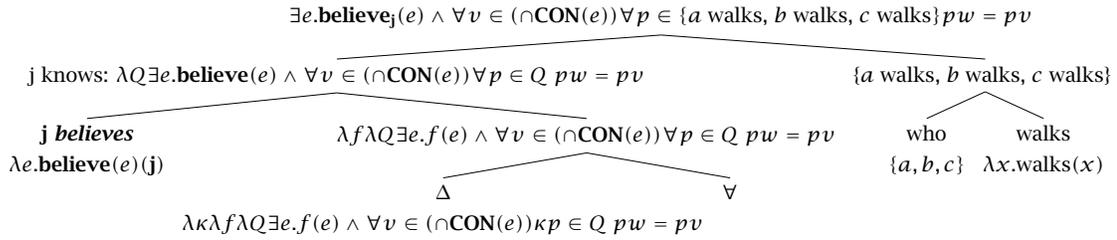


Figure 2: question-embedding under “know”

Now we’ve moved some pieces around, but how does this help? The answer should be quite clear by now: since now the answerhood condition is hardwired into the meaning of the verb in the form of a resolution layer – or really, in the form of a quantificational variable – it’s easy to have a verbal modifier attaching to this layer. I propose that’s exactly what the QVE adverbs are doing: *mostly* is basically just the vanilla quantifier **most**, in the QVE reading it attaches to the resolution layer  $\Delta$ , sutarating the  $\kappa$  variable. The full derivation of “John mostly knows who walks” is given in figure 3 below. The meaning we eventually derive

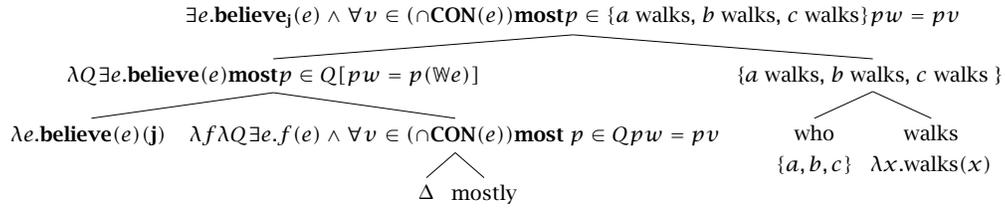


Figure 3: [[John knows who walks]] (QVE reading)

can be paraphrased as follows.

- (11) [[John mostly knows who walks]]:  
On most of the possible answers to the question “who walks”, John’s judgment is correct.

As desired, this is a reading that is sensitive to false answers: if John erroneously believes many people who doesn’t walk walks, the sentence will be judged as false. Of course, one may imagine a situation where the whole conversation goal is to finding the true answers and we may genuinely be agnostic about the mistakes the subject makes; this would be a case you want the Berman paraphrase – the quantified weakly exhaustive reading. The account here has the flexibility to account for

this reading as well, all we need to do is to change the world relation of  $\Delta$  from equivalence to entailment:

$$(12) \quad \Delta(\text{weak}) : \lambda \kappa_{a \rightarrow t} \lambda f_{v \rightarrow t} \lambda Q_{s \rightarrow t} . \exists e f(e) \wedge \forall v \in (\cap \text{CON}(e)) \kappa p \in Q \ p w \rightarrow p v$$

Accounting for the weakly exhaustive reading is never a problem, the whole point is how the “global view” of the strongly exhaustive reading can be quantified over. This account does the trick, by mediating on the answerhood conditions directly. However, given the fact that the adverb *mostly* is a modifier of the verb but not *just* the embedded question, the QVE data really seems to compel the non-trivial assumption we made in the beginning of this section: the answerhood condition is encoded by the verb.

## 5 Resolutive verbs

Now let’s go back to the other non-trivial assumption I made: *resolutive* verbs is a natural class of verbs. We’ve think of this in an intuitive way by comparing *know* and *believe*, but exactly what verbs can count as resolutive?

What matters to us is that those verbs contain a *non-trivial* resolution layer. Take *believe* as an example again, it is not resolutive, because its resolution would be really trivial: possibly some kind of equivalence relation between the subject’s belief worlds and their belief worlds – which comes out as a tautology. The same logic applies to rogative verbs like *wonder* and *ask*, they can’t be resolutive, because what would you resolve to? You form the questions you wonder about exactly when you start wondering about them. Therefore, I propose that *resolutive* verbs are attitude verbs whose *objects* can’t just hold in an initial mental state. That’s why the resolution layer is non-trivial, because these verbs inherently requires the mental objects of the initial state be resolved in another (set of) possible world(s) – the “evaluation world” we have been talking about. In the case of factive verbs, the evaluation world is the real world; for veridical verbs, the evaluation world is some kind of projected common ground (Anand & Hacquard (2014)); for some communication verbs like *discuss* and *agree on*, it could be the common ground of the participants of the event. Along this line, we have the following typology of embedding verbs in table 1.

Some final remarks on exactly how should we understand the “resolution” encoded by the verb. It worths to note that it’s more implausible to say this layer only arises when the complement of the verb is a question: in the case of embedding a proposition, the resolution part we propose for *know* is more likely a presupposition (or some backgrounded information), and for veridical verbs like *tell* is

resolutive	factive: <i>know, discover</i> non-factive: <i>tell, claim, agree on, demonstrate, discuss</i>
non-resolutive	rogative: <i>wonder, ask</i> anti-rogative: <i>believe, hope, wish</i>

Table 1: A typology of embedding verbs

more likely some sort of discourse move that can be cancelled (Anand & Hacquard (2014)). Farkas & Bruce (2010)'s idea is that these part of sentence meaning is the uptake intent of conversational acts, and should be formalized in terms of future states of the common ground. If their proposal is on the right track, the QVE data in this paper really raises another interesting issue: why a discourse move can be a target of quantification seems to be a target of quantification just when a question is embedded? The exploration of this question will be left to another occasion.

## References

- Anand, P., & Hacquard, V. (2014). Factivity, belief and discourse. *The art and craft of semantics: A festschrift for Irene Heim, 1*, 69-90.
- Beck, S., & Sharvit, Y. (2002). Pluralities of questions. *Journal of Semantics*, 19(2), 105-157.
- Berman, S. R. (1991). On the semantics and logical form of {it wh}-clauses.
- Farkas, D. F., & Bruce, K. B. (2010). On reacting to assertions and polar questions. *Journal of semantics*, 27(1), 81-118.
- Ginzburg, J. (1995a). Resolving questions, i. *Linguistics and philosophy*, 18(5), 459-527.
- Ginzburg, J. (1995b). Resolving questions, ii. *Linguistics and Philosophy*, 18(6), 567-609.
- Groenendijk, J., & Stokhof, M. (1989). Type-shifting rules and the semantics of interrogatives. In *Properties, types and meaning* (pp. 21-68). Springer.
- Groenendijk, J. A. G., Stokhof, M. J., et al. (1992). *Interrogatives and adverbs of quantification*. Institute for Logic, Language and Computation.
- Hacquard, V. (2010). On the event relativity of modal auxiliaries. *Natural language semantics*, 18(1), 79-114.
- Heim, I. (1994). Interrogative semantics and karttunen's semantics for know. In *Proceedings of iatl* (Vol. 1, pp. 128-144).
- Klinedinst, N., & Rothschild, D. (2011). Exhaustivity in questions with non-factives. *Semantics and Pragmatics*, 4(2), 1-23.
- Lahiri, U. (2002). *Questions and answers in embedded contexts*. Oxford University Press on Demand.