

Perhaps the Symmetry Problem solves itself

Matthijs Westera

University of Amsterdam / Universitat Pompeu Fabra

Manuscript, September 2017

Abstract. If the notion of relevance is symmetrical, pragmatic accounts of exhaustivity may fail to predict exhaustivity, yielding a contradiction or ignorance instead. This is the well-known Symmetry Problem. This paper presents a new pragmatic solution to the Symmetry Problem, discusses its compatibility with different types of pragmatic theories of exhaustivity, and compares it to existing pragmatic attempts to solve the Symmetry Problem. The paper demonstrates that the Symmetry Problem solves itself, within some but not all pragmatic accounts of exhaustivity, provided we construe relevance in terms of questions under discussion, and the latter as elements of a discourse strategy.¹

Keywords: Symmetry Problem, exhaustivity, relevance, question under discussion, discourse strategy.

1 Introduction

In this paper we adopt a pragmatic perspective on exhaustivity – though we will occasionally reflect, in passing, on the extent to which our arguments may have some bearing also on the grammatical approach (e.g., Chierchia et al. 2012). To our awareness all pragmatic accounts of exhaustivity assume that exhaustivity is the exclusion of propositions that are *relevant*, in some sense. To illustrate, consider the following typical example of an exhaustivity implication:

(1) A: Who (of John, Mary, Bill, Sue and Chris) was at the party?

B: John was there, and Bill was there.

(Exhaustivity: Mary, Sue and Chris weren't there.)

Here the proposition that Mary was there, for instance, is arguably relevant, and it is excluded by the exhaustivity implication. An important question for pragmatic accounts of exhaustivity is what should happen if relevance is *symmetrical*, e.g., if the set of relevant propositions is a partition or closed under negation. After all, one cannot exclude both a proposition and its negation, lest the utterance would imply a contradiction. This is the well-known *Symmetry Problem* (a term commonly attributed to MIT course notes by Heim and von Stechow). It was pointed out by Kroch (1972), subsequently discussed by Groenendijk

¹I am very grateful to Jeroen Groenendijk and Floris Roelofsen for valuable criticism on all iterations of the work presented here, as well as to two anonymous reviewers for SALT. Financial support from the Netherlands Organization for Scientific Research (NWO) is gratefully acknowledged.

and Stokhof (1984); Matsumoto (1995), among others, and more recently by, e.g., Katzir (2007), Fox and Katzir (2011), Block (2008) and Chierchia et al. (2012).

The symmetry problem can in principle be understood in two ways, the first of which appears to be more common:

- **Conceptual:** Relevance is necessarily symmetrical, in general or at least in cases where exhaustivity is implied.
- **Empirical:** Relevance is sometimes symmetrical in cases where exhaustivity is implied.

As a *conceptual* problem, the supposed symmetry of themes would presumably reflect some deep fact about what speakers are interested in or about the way in which they organize their conversational goals. In contrast, as a non-conceptual, *empirical* problem, it would really depend on the empirical facts whether the symmetry problem even exists. Example (1) may not make a particularly strong case for relevance being symmetrical. Potentially more problematic is the following, where A's initiative is explicitly symmetrical:

- (2) A: I need to know for all five people on this list (John, Mary, Bill, Sue, and Chris) whether they were present or absent.
B: John was present, and Bill was present.
A: Oh dear, only two, that's disappointing!

If B's response ends with falling intonation, A's subsequent conclusion seems to us warranted. This, then, may be a context in which the question introduced by A is undeniably symmetrical, yet in which B's utterance implies exhaustivity nevertheless – although of course this calls for a proper empirical evaluation.

In what follows, we aim to show first that the symmetry problem is probably not a conceptual problem. Subsequently, and more centrally, we aim to show that even if it is a conceptual problem, or if it exists as an empirical problem (i.e., (2)), the symmetry problem essentially solves itself. We will then discuss two existing approaches aimed at solving the symmetry problem within a pragmatic account, namely approaches based on conciseness/brevity and approaches based on Horn scales. We will show that neither of these is entirely satisfactory, although both share some insights with our solution. The paper ends with a conclusion and some discussion about broader implications.

2 Why the symmetry problem is probably not conceptual

Several authors seem to conceive of the symmetry problem as a conceptual problem, i.e., they conceive of relevance as being necessarily symmetrical (e.g., Block 2008; Fox and Katzir 2011; Chierchia et al. 2012). This view is not uncontroversial; it contrasts, for instance, with the view of Horn (1978) and Leech (1981) that speakers tend to be much more interested in what there is than in what there isn't – an instance of what Horn (2001) calls the *Asymmetry Thesis*. Still, the view that relevance would necessarily be symmetrical is quite widespread in the literature on exhaustivity, and we will discuss its motivation in some detail. We have found three general arguments for the symmetry of relevance:

1. that the interests of agents in a conversation would necessarily be symmetrical;

2. that the symmetry of relevance would follow from conceiving of relevance in terms of answerhood to a question;
3. that exhaustivity, being a conversational implicature, is part of what the speaker means, hence relevant.

We will discuss each in turn.

Argument 1. Symmetrical interests One possible argument in favor of the symmetry of relevance is that the interests of a speaker would necessarily be symmetrical, i.e., that speakers are interested in whether certain propositions are *true* or *false*. In support of this view, Chierchia et al. (2012) cite Carnap 1950 for a formal notion of relevance based on this idea, namely, based on a conception of agents as *hypothesis testers*; Fox and Katzir (2011) cite Lewis 1988 for a similar view. It seems to us that to this view of agents as hypothesis testers isn't necessarily appropriate for speakers in a conversation, and that a plausible alternative conception is available. We will briefly sketch such an alternative, and hope to make it somewhat plausible, though not in any sort of detail. This will be sufficient because, as we will explain shortly, whether a speaker's interests are necessarily symmetrical or not doesn't really matter.

Instead of conceiving of conversational agents as hypothesis testers – interested in the truth or falsity of a certain proposition – one could perhaps conceive of conversational agents as being interested not primarily in the truth or falsity of certain propositions, but in obtaining information that will either enable them to execute certain extra-conversational plans or require them to revise their plans (plans which will inevitably rely on certain defeasible assumptions). To illustrate, consider the following example:

- (3) (*B sees A confidently leaving the house without an umbrella, as usual.*)
- a. B: It's going to rain.
 - b. (?) B: It will stay dry.

Person A may be grateful for (3a) and adjust their plans by taking an umbrella. In contrast, (3b) is rather puzzling, because it is unhelpful, not worth sharing, and, we would say, “irrelevant”. Although one could of course disagree with our use of the English word “irrelevant” here, it seems to us hard to deny that there is some contrast between the two pieces of information in (3) that may be of pragmatic interest, and that one could safely assume a theory of conversational goals and hence a (technical) notion of relevance to be sensitive to.

However, as we mentioned, even if a speaker's interests are symmetrical (whether necessarily or occasionally), there may well be a reason for the speaker not to pursue or share all interesting pieces of information at the same time or by the same means. This possibility is entailed, for instance, by theories based on questions under discussion (QUDs; e.g., Roberts 1996): rather than pursuing all interesting pieces of information simultaneously, speakers organize them into QUDs (or topics, or themes) and pursue only some of those QUDs at any given time. According to this view, which we adopt, it is only relevance to the current QUD that matters for exhaustivity – and not all pieces of information that are in principle worth sharing need to be relevant in this technical sense. The organization of interesting propositions into QUDs is constrained in part by considerations of subject matter, of course, but also by considerations of overall orderliness, clarity and smoothness of the ensuing conversation. It may be beneficial, for instance, to split the complex QUD of *who ate what*

into the simpler sub-QUDs of *what John ate*, *what Mary ate*, and *what Fred ate* – a typical case of a discourse strategy (Roberts, 1996; Büring, 2003). Moreover, it has been noted that when such a strategy is (arguably) pursued, the relevant utterances (or utterance parts) will imply exhaustivity only relative to their respective sub-QUDs (Büring 2014; see also Wagner 2012). Hence, regardless of whether a speaker’s interests are (necessarily or occasionally) symmetrical, it is possible in principle for a speaker to organize a symmetrical set of interesting propositions into two distinct, asymmetrical QUDs, a positive and a negative one. Of course, linguists must not invoke such implicit QUD-maneuvers too freely; one would have to explain, primarily, why a speaker would prefer two asymmetrical QUDs over a single symmetrical one, and how an addressee might figure this out – we return to this important matter in section 3.

Argument 2. Questions as partitions Because sets of interesting propositions have come to be called questions under discussion, and because questions are often introduced by means of interrogative sentences, one would expect constraints on the semantic contents of interrogative sentences to reflect the structure of QUDs (though without conflating the two notions altogether). Indeed, Chierchia et al. (2012) motivate their assumption that relevance would necessarily be symmetrical by pointing to the semantics of interrogatives from Groenendijk and Stokhof 1984. This semantics assigns to interrogatives *partitions* on the set of possible worlds, i.e., sets of complete (exhaustive) answers.² If interrogative sentences express partitions, one would expect (relevance to) a QUD to exhibit a similar structure, and hence a similarly problematic symmetry.

However, this argument from partition theory is not entirely convincing. For one, Groenendijk and Stokhof themselves argue (p.528 and onwards) that it is necessary to pragmatically distinguish positive partial answers from negative partial answers, a matter for which they note partitions are too coarse (unlike, for instance, the *question abstracts* in their account). But more importantly, alternative semantic theories of interrogatives exist that do not exhibit symmetry, e.g., Hamblin 1973, or more recently Ciardelli et al. 2015. The latter – inquisitive semantics – even satisfies the logico-philosophical desiderata which Groenendijk and Stokhof thought committed them to partitions, foremost, that on disjunction and conjunction of interrogatives should correspond to union and intersection in the usual way (for discussion see Ciardelli 2014). Hence, rather than invoking partition semantics as an argument for a symmetrical notion of relevance, it seems equally possible to invoke the symmetry problem as an argument against partition semantics (or, alternatively, against a direct correspondence between the semantics of interrogatives and the structure of QUDs).

Argument 3. Exhaustivity is meant, hence relevant The third argument that we have found in the literature – but we have unfortunately lost track of its exact source – starts from the assumption that, in the relevant examples, exhaustivity is an *implicature* and hence part of what the speaker *means* – an assumption which we share. Now, if indeed both the explicit assertion and the exhaustivity are part of what the speaker means, then both parts are subject to the Gricean maxims, including Relation, hence both the positive and the negative information must be relevant. Accordingly, so the argument goes, relevance must be symmetrical. However, we think that this conclusion is not warranted if relevance is understood, again, in terms of QUDs, for the following reason.

²Although partitions are not strictly closed under negation, they do exhibit the problematic type of symmetry, and, besides, closure under negation can be achieved by also taking all unions of partition cells into account, i.e., all partial answers.

The information provided by an utterance is not a monolithic blob, but structured in a certain way. In particular, the primary, explicit assertion (or Grice’s (1989) “what is said”) and conversational implicature constitute different intents (Neale 1992; Bach 2006; and the same holds for *conventional* implicature, e.g., Potts 2005; Simons et al. 2010), and as such they may well be aimed at distinct QUDs.³ Indeed, they *must* be aimed at distinct QUDs according to common definitions of the maxim of Quantity (e.g., Harnish 1976; Gamut 1991; Schulz and Van Rooij 2006; Spector 2007), which we think are basically correct: no two distinct intents can comply with Quantity relative to the same QUD, because at least one of the two intents (and in our examples both) would have to be underinformative. So although both are part of what is meant, and hence both must be relevant in the broad sense of being worth sharing, they can and must be relevant to distinct QUDs. Crucially, neither of these QUDs needs to be symmetrical in principle, and in addition the two QUD need not be equally pertinent/central/important (just as conventional implicatures would be aimed at secondary, ancillary QUDs, e.g., (Potts, 2005)). Thus, while argument 3. may be an argument for some degree of symmetry in a speaker’s interests (in the relevant examples), it is not an argument for symmetrical QUDs, i.e., for the symmetry of a more narrow notion of relevance that matters for exhaustivity. Of course the foregoing leaves unexplained, just like our response to argument 1, why a speaker would address two asymmetrical QUDs rather than a single symmetrical QUD, and how an addressee (or a linguist) might detect such a maneuver. These important matters will be dealt with in the next section.

Summing up, even if speaker interests are in fact symmetrical – whether in general, which we called into question, or occasionally – a speaker may in principle subdivide this symmetrical set of interesting propositions into two asymmetrical QUDs, and address one by means of an assertion and the other by means of an exhaustivity implicature. This follows from the more general fact that within the QUD-based approach to pragmatics, speaker interests are only one factor in the organization of conversational goals. Of course we have yet to identify the other crucial factors for cases involving exhaustivity – and until we do so, the suggested two-QUD analysis of exhaustivity may appear terribly ad hoc.

3 How the symmetry problem solves itself

Let us set aside the question of whether or not speaker interests are necessarily symmetrical – i.e., the question of whether the symmetry problem is a conceptual problem – by concentrating on example (2), repeated here, where the speakers’ interests are arguably symmetrical regardless:

- (2) A: I need to know for all five people on this list (John, Mary, Bill, Sue, and Chris) whether they were present or absent.
 B: John was present, and Bill was present.

³Exhaustivity is commonly called “implicature” also in the grammatical approach, but proposals in this strand seem to derive it as part of the primary assertion, which would in fact commit one to a symmetrical QUD for the relevant utterances. Since the grammatical approach does not necessarily seek to derive exhaustivity from something like a QUD, this symmetry is not necessarily problematic. Nevertheless, we think that a multi-intent, multi-QUD analysis would be possible and worth exploring also within the grammatical approach, e.g., by treating exhaustivity as a *conventional* implicature and computing it through an appropriate multi-dimensional semantics (e.g., Potts 2005; Gutzmann 2015). But we must leave this for another occasion.

A: Wow, only two, what a disappointment!

For ease of speaking, let us suppose that A’s final response in (2) is indeed conceivable, i.e., that (with the right intonation) B’s preceding answer can imply exhaustivity; that is, let us suppose that the symmetry problem as an empirical problem really exists. Whether this is true is of course an important question that calls for experimental work, but this must be left for another occasion. At the end of this section we will briefly reflect on the possibility that our judgment regarding (2) may not be representative, i.e., the possibility that B’s response would not in fact imply exhaustivity.

At this point it is necessary to be a bit more explicit about pragmatic theories of exhaustivity. Let us assume that A’s initiative in (2) introduces the following QUD, where Px denotes the proposition that x is at the party, \bar{X} denotes set complementation, and X^\cap denotes closure under intersection:

$$\text{QUD} = \{Pj, Pm, Pb, Ps, Pc, \bar{Pj}, \bar{Pm}, \bar{Pb}, \bar{Ps}, \bar{Pc}\}^\cap$$

It is useful to distinguish theories of exhaustivity based on what they (try to) predict if we take B’s answer in (2) to comply with the conversational maxims relative to this symmetrical QUD. Theories may (try to) predict either:

- I. **contradiction:** the speaker’s epistemic state is inconsistent (due to excluding both Pm, Ps and Pc , and \bar{Pm}, \bar{Ps} and \bar{Pc});
- II. **ignorance:** the speaker is unsure whether Mary, Sue and/or Chris were present (due to excluding neither Pm, Ps and Pc , nor \bar{Pm}, \bar{Ps} and \bar{Pc});
- III. **exhaustivity:** the speaker believes that Mary, Sue and Chris were absent (due to excluding only Pm, Ps and Pc).

Of type II. is the plain vanilla recipe for exhaustivity, based on the Gricean maxim of Quantity and a contextual default assumption of opinionatedness (e.g., Soames 1982; for discussion see Sauerland 2004; Geurts 2011): compliance with Quantity implies the speaker’s inability to assert either of the two symmetrical alternatives, i.e., ignorance, and this contradicts the purported opinionatedness assumption which, it being a mere default, is simply dropped (or weakened; Schulz and Van Rooij 2006). Of type III. are common refinements of the standard recipe that try to break the symmetry, for instance, by assuming that negative propositions are significantly more complex to convey than positive ones, and/or that relevance would be filtered by intrinsically asymmetrical, lexical scales. Theories of type I., finally, are less common – we are aware only of Westera (2014), where opinionatedness is not a contextual default assumption but entailed by compliance with the maxim of Relation, and Westera (2017), where opinionatedness is bypassed altogether by deriving exhaustivity from a set of maxims that govern attention sharing.⁴

Superficially it might seem that the Symmetry Problem would be the most serious threat to theories of type I., which predict a contradiction, less so for type II., and even less so for theories of type III., which after all claim to solve it. But we aim to show that the converse is true: theories of type III. are in fact problematic, whereas theories of type I. enable a satisfactory solution, with type II. somewhere in between. This is because, as announced by

⁴In a nutshell, while B’s response in (2) may well comply with the Gricean, information-governing maxims relative to the symmetrical QUD, it violates Westera’s (2017) attentional maxim of Quantity, according to which a speaker should draw attention, i.e., mention, every relevant possibility. After all, if both Mary’s presence and Mary’s absence are relevant, speaker B should consider at least one of them possible, but she didn’t mention either.

the title of this paper, the Symmetry Problem *solves itself* – and, paradoxically, the greater the problem, the better it can solve itself. In what follows we will adopt the perspective of a type I. theory, which predicts a contradiction, and present the solution this offers. Afterwards we will briefly discuss the extent to which a similar solution may be available for type II. theories. In each case the internal mechanics of these theories will not play a role and will be left implicit – what matters is only their output. Theories of type III., based on conciseness and/or scales, will be discussed in some detail separately in section 4.

Within theories of type I., the assumption that B’s response in example (2) would comply with the maxims relative to the symmetrical QUD leads to a contradiction. Put differently: B’s response cannot compliantly address that QUD. Although speakers may in principle violate maxims, namely in case of a clash, as Grice (1989) noted they must not do so silently, lest they be liable to mislead. We assume, following Westera 2013, that maxim violations are often signaled prosodically, by a final rising contour (or in written text by, e.g., “...”). In the absence of such cues, as in example (2), only one conclusion is possible: contrary perhaps to appearances, B’s answer must be aimed at a different QUD, i.e., different from the symmetrical QUD introduced by A. Note that this is not an assumption invoked solely to solve the Symmetry Problem, which would be ad hoc; rather, regardless of how one might try to solve the Symmetry Problem, a QUD-shift is simply *predicted* by theories of type I.⁵ What remains of the Symmetry Problem for theories of type I. is not really a problem, but a set of research questions:

- (i) Which QUD is (or which QUDs are) addressed by speaker B in (2), if not the symmetrical QUD introduced by speaker A?
- (ii) Why was this a rational choice for B?
- (iii) How can an addressee (e.g., speaker A) figure this out, accommodate the new QUDs and compute the right inferences?

We will address each in turn.

Question (i): Which QUDs are addressed by B? As announced in the previous section, we propose that for some reason speaker B in (2) decided to split the prior symmetrical QUD into two asymmetrical QUDs, which we shall denote by QUD^+ and QUD^- :

$$\text{QUD}^+ = \{Pj, Pm, Pb, Ps, Pc\}^\cap \quad \text{QUD}^- = \{\overline{Pj}, \overline{Pm}, \overline{Pb}, \overline{Ps}, \overline{Pc}\}^\cap$$

Let us assume, we believe uncontroversially, that an utterance that addresses multiple QUDs should convey an appropriate communicative intent for each QUD, i.e., an intent which complies with the maxims relative to that QUD. This explains why B’s response in (2) would be fine with the assumed QUDs, in the following way:

1. B’s primary (asserted/explicit) intent is that John and Bill were at the party, which can safely comply with the maxims relative to QUD^+ (under any type of pragmatic theory of exhaustivity);
2. because QUD^+ is asymmetrical, compliance with the maxims of the primary intent implies exhaustivity as usual, i.e., that according to the speaker Mary, Sue and Chris were absent (under any type of pragmatic theory of exhaustivity);

⁵In comparison, the assumption of a QUD-shift would be considerably more ad hoc within theories of type II. or III., according to which the B’s response in (2) may well comply with the maxims relative to a symmetrical QUD.

3. the exhaustivity implication in turn enables the clear communication of a secondary intent, i.e., an conversational implicature, namely that Mary, Sue and Chris were absent;⁶
4. the secondary intent can safely complies with the maxims relative to QUD⁻ (under any type of pragmatic theory of exhaustivity).

That is, instead of addressing “who was present and who was absent?”, for some reason speaker B decided to address only the positive half explicitly, enabling her to address the negative half implicitly by means of an exhaustivity implicature.

Question (ii): Why was this QUD-split rational? Splitting a QUD into two is an ordinary *discourse strategy* (Roberts, 1996), and one which in this particular situation offers a substantial benefit over simply addressing the original, symmetrical QUD. The benefit is that addressing an asymmetrical QUD enables an exhaustivity implicature, unlike the original symmetrical QUD – that’s the Symmetry Problem – and that the exhaustivity implicature enables the speaker to address half of the original QUD, i.e., the other asymmetrical QUD, implicitly, which greatly benefits conciseness. In this regard the symmetry problem solves itself: an asymmetrical QUD is favored precisely because the symmetrical QUD prevents an exhaustivity implicature. Now, this doesn’t mean that this kind of QUD-split is always appropriate. For instance, if speaker A is ticking boxes on a checklist of individuals, it might be better for B to address the entire original QUD explicitly, and in the precise order of the checklist:

- (4) B: John was present, Mary absent, Bill present, Sue absent, and Chris absent.

Moreover, addressing half of the original QUD implicitly may not be a good idea in cases where the domain of relevant individuals is not entirely clear, which would compromise the clear communication of the exhaustivity implicature. But in other circumstances B’s decision in (2), to split the QUD, seems to be perfectly rational.

Note that speaker B could also have chosen to address the negative QUD explicitly, explicating who was *absent* and implicating that the others were *present*:

- (5) A: I need to know for all five people on this list (John, Mary, Bill, Sue, and Chris) whether they were present or absent.
 B: Mary was absent, Sue was absent and Chris was absent.
 A: That’s a lot of absences. Well, it’s good to learn that John and Bill were there!

In this particular scenario, where only two people were present, implicating who was present offers a smaller conciseness benefit than implicating who was absent. In general, the more people were present (and also the greater the domain of A’s inquiry) the greater the conciseness benefit of explicating only who was absent. But other factors will also play a role, such as which of the two properties, being present or being absent, is the most salient in the broader context, e.g., (5) seems particularly natural if B normally takes attendance by writing down only the names of those who are absent. Other factors that may play a role are, for instance, which of the two predicates is the most readily lexically accessible,

⁶The distinction between implication and implicature is important (Bach, 2006): what is implied is not necessarily implicated (meant), and what is implicated is not necessarily implied to be true (but, typically, implied to be held true by the speaker).

and which of the individuals' names B is more likely to mispronounce. But for present purposes we set such complications aside, because the main conciseness benefit is more general: relying on conversational implicature for conveying part of the answer implicitly benefits conciseness regardless of the particular lexical entries involved. In this regard our proposal is crucially different from existing conciseness-based approaches to the Symmetry Problem, as we will see in section 4.

While the conciseness benefit alone may explain why B chose to split the prior QUD into two halves, it does not explain why it should be split into a positive QUD and a negative QUD, i.e., QUD^+ and QUD^- , rather than, e.g.:

$$\text{QUD}_1 = \{Pj, \overline{Pm}, Pb, \overline{Ps}, Pc\}^\cap \quad \text{QUD}_2 = \{\overline{Pj}, Pm, \overline{Pb}, Ps, \overline{Pc}\}^\cap$$

After all, this split could have offered, depending on who was actually at the party, a similar conciseness benefit as the split into QUD^+ and QUD^- (though with different predicted exhaustivity implicatures). One reason why the above split may be dispreferred is that the resulting QUDs are more complex: the propositions in QUD^+ and QUD^- vary only along a single dimension, i.e., the individual, whereas the propositions in QUD_1 and QUD_2 vary along two dimensions, i.e., the individual, and whether they were absent or present – and they vary in a rather unpredictable way, including some but not all combinations of individual and absence/presence. This added complexity would compromise clarity: an addressee may not be able to figure out which of many possible asymmetrical-but-mixed QUDs the speaker may be addressing (also if we take prosodic focus into account, see below). We take this to explain why the QUD-split must be as assumed, i.e., into QUD^+ and QUD^- .

Question (iii): How could an addressee figure this out? The main answer to this question is that, according to theories of type I., B's response in (2) *must* be aimed at a different QUD. After all, it would have violated a maxim relative to the symmetrical QUD – and addressees should recognize this. In this regard, too, the symmetry problem solves itself. Which different QUD(s) speaker B may be addressing is constrained, in turn, by the notion of discourse strategy: it must be some combination of QUDs that together cover the original one, and we have already explained why the assumed split into a positive and a negative QUD is favored.

Besides these general pragmatic considerations, and independently of one's precise theory of exhaustivity, in spoken language addressees may also rely on *prosodic focus* for identifying the QUD. According to our judgments, for B's response in (2) to imply exhaustivity, it should have a pitch accent on the individuals' names but not on the predicates (just as the standard focus-congruence pattern for the original example, (1)). Let us assume that accent placement reflects, through focus structure in the usual manner (e.g., Rooth 1985; Beaver and Clark 2009), only the primary QUD, i.e., the QUD that is explicitly addressed. The focus structure of (2) will then help an addressee to understand that the primary QUD of B's response is the asymmetrical, positive one. In contrast, if B had been addressing the symmetrical QUD, or a strange mixture like QUD_1 or QUD_2 above, B should have used either broad focus (i.e., on the entire sentence, which would normally entail an accent at least on the predicates) or multiple foci, i.e., both the individuals' names and the predicates. To illustrate, example (4) does address the symmetrical QUD, and indeed an intonation contour with accents on both the names and the predicates seems to us the most natural.

Summing up, theories of exhaustivity of type I., which predict a contradiction relative to a symmetrical QUD, enable a solution to the Symmetry Problem precisely because this

contradiction entails a QUD-shift. To explain the QUD-shift we invoked ordinary discourse strategies, and noted that splitting a symmetrical QUD into two asymmetrical QUDs offers a conciseness benefit by virtue of enabling exhaustivity implicature. In contrast, theories of type II., which merely predict ignorance relative to a symmetrical QUD, do not force a QUD-shift: according to these theories B’s response may well comply with the maxims relative to a symmetrical QUD, and it is only the default assumption of opinionatedness that needs to be given up. As such, we think that the proposed QUD-manuever is at best *compatible* with approaches of type II., and only if prosodic focus is available to sufficiently guide the addressee (i.e., as an answer to question (iii)).

Lastly, as we announced, let us briefly consider the possibility that our judgments regarding (2) (and also (5)) are not in fact empirically representative, i.e., that B’s response would actually fail to imply exhaustivity in a situation where speaker interests are explicitly symmetrical. For instance, if instead B’s response would turn out to be judged infelicitous, that could be an argument for theories of type I., but against the type of discourse strategy on which we relied – we think that such a finding would pose a significant challenge to the broader approach based on QUDs. If instead B’s response is as felicitous as we think it is, but it would turn out (even with falling intonation) to imply ignorance rather than exhaustivity, this could be an argument for theories of exhaustivity of type II., and against theories of types I. and III. Lastly, if instead B’s response would turn out to imply neither exhaustivity nor ignorance, we believe that any pragmatic account of exhaustivity would face a serious challenge.⁷ If, on the other hand, our judgments about (2) are representative, this would constitute evidence for theories of type I., and against theories of type II. In the next section we discuss theories of type III.

4 Existing approaches to the symmetry problem

Type III. theories of exhaustivity, recall, try to predict exhaustivity despite a symmetrical QUD. Strictly speaking, our proposal in the previous section, on the basis of a type I. theory, achieves something similar, but only by first replacing the symmetrical QUD by an asymmetrical one. As we will see, type III. theories are substantially different from our proposal, although there are also some commonalities. We will discuss two strands of approaches: those based on (Horn) scales, and those based on considerations of conciseness.

4.1 Approaches based on scales

Scales of the sort introduced by Horn (1972) were invoked by Gazdar (1979) as a potential solution to the Symmetry Problem: lexical entries would be associated with certain intrinsically asymmetrical scales of alternatives, which could serve as filters on symmetrical QUDs. The Symmetry Problem in cases like (2) is then solved, superficially at least, because although “Mary is absent” is relevant, it is not a scalar alternative to “John is present” and as such does not enter into the derivation of exhaustivity. However, Russell (2006) points

⁷One might in that case try to save these theories by assuming that B implicitly restricted the QUD’s domain, say, to those individuals that B finds *particularly* relevant (perhaps implying that B didn’t quite agree with the scope of A’s inquiry). To close this escape hatch, a better test case might be to have speaker B explicate the domain as well:

- (1) B: Of those five individuals you mention, John was present, and Bill was present

out that scales aren't really a solution to the Symmetry Problem unless one explains why scales are the way they are, and why they should function as filters on QUDs prior to the computation of exhaustivity; and Geurts (2011) notes that there is only very little explicit reflection on what scales are supposed to be. Let us go through what we believe are the four main conceptions of scales in the literature, and compare them to our proposal.

The first option is to conceive of scales, following Lassiter (2010), Geurts (2011) and, we think, Horn (1972), as indirect representations of what is *typically* relevant given that a certain lexical expression is used. Although knowledge of what is typically relevant is of course important in communication, one would expect it to play a role only in cases where the audience is left to guess what the QUD is. Hence, one would not expect it to play a (large) role in cases like (2), where there is an explicit preceding question, which together with discourse strategies, prosodic focus and, given a theory of type I., the necessity of a QUD-shift, arguably suffices for an addressee to figure out what the new QUDs are. Given this, knowledge of what is typically relevant may at best quicken the audience's comprehension; it does not play a substantive role in the explanation. In fact, our explanation for why speakers would avoid symmetrical QUDs may even be part of the reason why QUDs would tend to be asymmetrical to begin with, i.e., why scales (as representations of typical QUDs) would be asymmetrical.

The second option is to conceive of scales, following, we believe, Hirschberg (1985) and Levinson (1983), as representations not of what is typically relevant but of what is *actually* relevant for a given utterance (in this role scales are also called "Hirschberg scales" or "ad hoc scales"; Huang 2014). Given this conception of scales, to say that in the relevant examples exhaustivity is computed with regard to a scale rather than a symmetrical QUD is just to say, in a more cryptical way, that the speaker shifted from the symmetrical QUD to an asymmetrical QUD, though without explaining why. While essentially in line with our proposal, we are unaware of any explanation for why such a QUD-shift would be rational and how an addressee might figure it out, and without such an explanation the reliance on scales (conceived of as representations of QUDs) is ad hoc. Our main contribution, in this regard, is that in the previous section we did provide such an explanation.

A third option is to conceive of scales as representations not just of what is actually/typically relevant, but of what is actually/typically relevant *and sufficiently concise/simple to express* (e.g., McCawley 1978; Atlas and Levinson 1981; Horn 1984). This may have some explanatory potential, and we will discuss this strand separately and in more detail in the next subsection.

A fourth option is to conceive of scales as lexical/grammatical devices that have nothing necessarily to do with QUDs or relevance. This type of interpretation has been suggested within the grammatical approach (e.g., Chierchia et al. 2012), which as announced we will not discuss in detail in the current paper. In the grammatical approach, scales are a powerful and perhaps necessary tool from a descriptive point of view, but their explanatory potential is currently unclear.

4.2 Approaches based on conciseness

Approaches based on conciseness would attempt to solve the symmetry problem in cases like (2) by assuming that "John was absent/wasn't there" is a more complex expression than "John was present/was there", and likewise for the other individuals. This would provide speakers with an excuse for omitting Mary's absence but not for omitting Mary's presence, thus breaking the problematic symmetry. For a recent proposal in this vein see Lassiter 2010.

Several authors have tried to define an appropriate notion of conciseness/complexity, for instance in terms of number of syllables (McCawley, 1978) or degree of lexicalization (Atlas and Levinson, 1981; Horn, 1984).⁸ Setting aside the interesting question of whether “absent” is really more complex than “present” (which an anonymous reviewer suggests may well be the case), what these proposals crucially need to assume is that the purported difference in complexity would be sufficiently large for it to matter. That is, the difference should be sufficiently large to provide speakers with an excuse for not mentioning certain relevant propositions, *and* for this omission to not cause any confusion among the addressees. This is something which we are hesitant to accept (in line with Carston 2005). We think it would require independent confirmation before this type of conciseness-based approach to the symmetry problem could be deemed satisfactory.

But even (or especially) if it can be shown that “absent” is significantly more complex than “present”, this kind of approach appears not to be on the right track. Consider again the mirror image of (2), i.e., (5), repeated here:

- (5) A: I need to know for all five people on this list (John, Mary, Bill, Sue, and Chris) whether they were present or absent.
B: Mary was absent, Sue was absent and Chris was absent.
A: That’s a lot of absences. Well, it’s good to learn that John and Bill were there!

Let us suppose again that our judgment that B’s response implies exhaustivity is representative of the empirical facts. In order for conciseness considerations to be able to explain this, conciseness would have to provide B with a reason for omitting “John was present” but not for omitting “John was absent”, which is the exact converse of what was needed for (2). Therefore, example (5) shows that a conciseness-based solution to the symmetry problem that feeds only on intrinsic properties of particular lexical entries will be inadequate; rather, there must be a contextual parameter of, say, “mentionworthiness”, that has nothing necessarily to do with intrinsic conciseness or complexity. A similar criticism is voiced by Matsumoto (1995), based on cases where a simple expression and a more complex expression are used together, e.g.:

- (6) B: It was warm today, and a little bit more than warm yesterday.

Matsumoto observes that the utterance implies that it was not a little bit more than warm today, despite this being expressible only by a more complex utterance. In response, Katzir (2007) proposes that sometimes complex expressions can be used in spite of their complexity, and that one can find out whether complex expressions can be used by checking whether the utterance itself contains such a complex expression somewhere. (Lassiter (2010) presents a similar view in defense of the conciseness-based approach to the Symmetry Problem.) This is of course true, but it doesn’t explain why a particular context would be such that the more

⁸ Katzir (2007) tries to filter something like relevance in terms of a measure of grammatical complexity, in a way that would superficially seem to belong in the same strand as the aforementioned approaches. However, Katzir does not intend this to be part of a pragmatic explanation, and indeed it is difficult to see how it could be. Katzir’s measure of grammatical complexity is defined in terms of whether certain permissible substitutions enable one to transform one sentence into another. A consequence of this is that which of two expressions counts as more complex in Katzir’s sense can depend on which expression was actually uttered. Although we can see the appeal of this proposal within the grammatical approach to exhaustivity, we do not see how it could follow from a global pragmatic preference for conciseness. For this reason we leave a discussion of Katzir’s proposal for another occasion, along with, as announced, a discussion of the grammatical approach more generally.

complex expression could be used to begin with. What it shows is that conciseness-based approaches must invoke a contextual parameter of “mentionworthiness” that is at least in part independent of considerations of intrinsic conciseness or complexity. It is worth noting that neither author considers an example like (5), which we think makes this even more evident.

Once the need for a contextual “mentionworthiness” parameter is acknowledged, which given example (5) cannot have anything necessarily to do with intrinsic lexical conciseness or complexity, let us just call the set of propositions that are worth mentioning a “QUD” and get rid of whatever symmetrical notion of relevance was used before (we can always obtain it by closing the notion of mentionworthiness or QUD under negation, should we find a need for it). The resulting picture is essentially what we motivated in section 3, by explaining why a speaker would choose to address an asymmetrical QUD despite symmetrical interests – and recall that intrinsic lexical conciseness or complexity did not play a role in our explanation. Rather, our explanation relied on the obvious conciseness benefit of conversationally implicating part of the answer, which obtains in (2) and (5) alike. Another important difference is that, in the solution we proposed, considerations of conciseness are not strictly necessary for an audience to be able to identify the exhaustivity implicature: conciseness may help explain *why* a speaker chose to address an asymmetrical QUD, but *that* the speaker did so will be evident regardless, from prosodic focus and (in a type I. theory) from the fact that the utterance would have violated a maxim otherwise. In contrast, according to previous conciseness-based approaches, the audience would not be able to understand the exhaustivity implicature except through taking conciseness into account. That such relatively tiny conciseness differences would play such a central role does not seem plausible, as Carston (2005) notes.

5 Conclusion and discussion

We have argued that the Symmetry Problem is not a conceptual problem: even if a speaker’s interests are symmetrical – whether in general, which we questioned, or occasionally – it may be rational for the speaker to organize the propositions of interest into asymmetrical QUDs – and the latter are what matters for exhaustivity. Solving what remained of the Symmetry Problem amounted to identifying such reasons, and here the Symmetry Problem turned out to solve itself: speakers may choose to split a symmetrical QUD into two halves precisely because a symmetrical QUD prevents them from conveying part of the answer implicitly, i.e., as an exhaustivity implicature.

Within the QUD-based approach we adopted, our solution in terms of a discourse strategy is forced upon us by accounts of exhaustivity that predict a contradiction (maxim violation) relative to a symmetrical QUD (i.e., type I. theories), while being compatible at best (namely, given prosodic focus) with accounts that predict ignorance instead (i.e., type II.). As for existing approaches that have tried to predict exhaustivity in spite of a symmetrical QUD (i.e., type III.), we discussed those based on scales and those based on conciseness. With (one interpretation of) the former our proposal shares that exhaustivity is computed relative to an asymmetrical QUD, and our new contribution is an explanation of why this would be so. With the latter our proposal shares that considerations of conciseness have some role to play, though with important differences in the kind of role: our explanation relies only on the general fact that conversational implicature benefits conciseness, regardless of the particular lexical entries involved. Altogether, if our proposal is on the right track, this

shows, paradoxically, that the best way to solve the Symmetry Problem may be to embrace it, i.e., to adopt a set of conversational maxims that imply a contradiction relative to a symmetrical QUD. This provides some independent support for pragmatic accounts (of type I.) that try to derive exhaustivity directly from the maxims, without an opinionatedness assumption (Westera, 2014, 2017).

As we announced in the introduction, we did not discuss the grammatical approach to exhaustivity, except briefly, in passing. Chierchia et al. (2012) list the Symmetry Problem as an argument (one of several) against pragmatic theories of exhaustivity, and in favor of the grammatical approach. We hope to have shown that this argument falls short: a pragmatic solution is available. Within the grammatical approach itself, symmetry is sometimes relied upon in order to block certain undesirable exhaustivity inferences. For instance, in order to block the “not all”-inference of the disjunction “some or all”, local exhaustification of the first disjunct would render the two disjuncts mutually exclusive, which because of their symmetry would block subsequent global exhaustification (e.g., Chierchia et al. 2009; Katzir and Singh 2013). Depending on how the sets of alternatives in the grammatical approach relate to something like relevance or QUDs – and to our awareness there is no consensus in this regard – our arguments in the current paper may have some bearing on the grammatical approach as well. We leave an exploration of this relation to future work.

Zooming out a little, our proposal highlights an important *division of pragmatic labor*, namely between choosing certain QUDs to address and selecting appropriate communicative intents and expressions for doing so. More generally, this is a division between choosing/organizing one’s goals and selecting the appropriate means to achieve them. Existing conciseness-based approaches to the Symmetry Problem have concentrated on the means, by keeping the symmetrical QUD in place and comparing the conciseness benefits only of different ways of addressing that QUD. Our proposal, in contrast, considered the conciseness benefits of a maneuver at the level of QUDs, i.e., of a certain discourse strategy, and this is what made it both explanatory and more successful, we believe, at dealing with certain challenging examples. We think it is essential for the field to keep this division of pragmatic labor in mind, and to explicate and carefully motivate assumptions at both levels.

References

- Atlas, J. D. and S. C. Levinson (1981). It-clefts, informativeness and logical form: radical pragmatics (revised standard version). In P. Cole (Ed.), *Radical pragmatics*, pp. 1–62. New York: Academic Press.
- Bach, K. (2006). The top 10 misconceptions about implicature. In *Drawing the boundaries of meaning: Neo-Gricean studies in pragmatics and semantics in honor of Laurence R. Horn*, pp. 21–30. John Benjamins Publishing Company.
- Beaver, D. and B. Clark (2009). *Sense and Sensitivity: How Focus Determines Meaning*. Number 12 in Explorations in Semantics. John Wiley & Sons.
- Block, E. (2008). Is the symmetry problem really a problem? Unpublished manuscript; retrieved from <http://web.eecs.umich.edu/~rthomaso/lpw08/block.pdf>.
- Büring, D. (2003). On d-trees, beans, and accents. *Linguistics and Philosophy* 26, 511–545.
- Büring, D. (2014). (Contrastive) topic. In C. Féry and S. Ishihara (Eds.), *The Oxford Handbook of Information Structure*. Oxford University Press. Online version.
- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.

- Carston, R. (2005). Relevance Theory, Grice and the neo-Griceans: a response to Laurence Horn's 'current issues in neo-Gricean pragmatics'. *Intercultural Pragmatics* 2, 303–319.
- Chierchia, G., D. Fox, and B. Spector (2009). Hurford's constraint and the theory of scalar implicatures. *Presuppositions and implicatures* 60, 47–62.
- Chierchia, G., D. Fox, and B. Spector (2012). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In C. Maienborn, P. Portner, and K. von Stechow (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, Volume 2, pp. 2297–2332. Mouton de Gruyter.
- Ciardelli, I. (2014). Question meanings = resolution conditions. Slides presented at the seventh Semantics and Philosophy in Europe Colloquium, Berlin; retrieved from <http://www.ivanociardelli.altervista.org/wp-content/uploads/2014/02/SPE-presented.pdf>.
- Ciardelli, I., J. Groenendijk, and F. Roelofsen (2015). On the semantics and logic of declaratives and interrogatives. *Synthese* 192(6), 1689–1728.
- Fox, D. and R. Katzir (2011). On the characterization of alternatives. *Natural Language Semantics* 19(1), 87–107.
- Gamut, L. T. F. (1991). *Language, Logic and Meaning (vols. 1 and 2)*. Chicago University Press.
- Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition, and Logical Form*. New York: Academic Press.
- Geurts, B. (2011). *Quantity Implicatures*. Cambridge University Press.
- Grice, H. (1989). *Studies in the Way of Words*. Harvard University Press.
- Groenendijk, J. and M. Stokhof (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*. Ph. D. thesis, University of Amsterdam.
- Gutzmann, D. (2015). *Use-Conditional Meaning: Studies in Multidimensional Semantics*. Oxford University Press. Oxford Studies in Semantics and Pragmatics 6.
- Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language* 10, 41–53.
- Harnish, R. (1976). Logical form and implicature. In T. G. Bever, J. J. Katz, and D. T. Langendoen (Eds.), *An Integrated Theory of Linguistic Ability*. New York: Crowel.
- Hirschberg, J. (1985). *A Theory of Scalar Implicature*. Ph. D. thesis, University of Pennsylvania.
- Horn, L. R. (1972). *On the Semantic Properties of Logical Operators in English*. Ph. D. thesis, University of California Los Angeles.
- Horn, L. R. (1978). Lexical incorporation, implicature, and the least effort hypothesis. In D. Farkas, W. M. Jacobsen, and K. W. Todrys (Eds.), *CLS: Papers from the Parasession on the Lexicon*, pp. 196–209. Chicago Linguistic Society.
- Horn, L. R. (1984). Towards a new taxonomy of pragmatic inference: Q-based and R-based implicatures. In D. Schiffrin (Ed.), *Meaning, Form, and Use in Context*, pp. 11–42. Georgetown University Press.
- Horn, L. R. (2001). *A Natural History of Negation*. David Hume Series on Philosophy and Cognitive Science Reissues. Stanford: CSLI Publications. First published in 1989.
- Huang, Y. (2014). *Pragmatics*. Oxford Textbooks in Linguistics. Oxford University Press.
- Katzir, R. (2007). Structurally-defined alternatives. *Linguistics and Philosophy* 30(6), 669–690.
- Katzir, R. and R. Singh (2013). Hurford disjunctions: embedded exhaustification and structural economy. In U. Etxeberria, A. Fălauş, A. Irurtzun, and B. Leferman (Eds.), *Proceedings of Sinn und Bedeutung 18*.
- Kroch, A. (1972). Lexical and inferred meanings for some time adverbs. *Quarterly Progress Reports of the Research Laboratory of Electronics* 104, 260–267.
- Lassiter, D. (2010). Why symmetry is not a problem for a gricean theory of scalar implicature.

- Presented at Utterance Interpretation and Cognitive Models 3; retrieved from <http://web.stanford.edu/~danlass/Lassiter-SI-UICM.pdf>.
- Leech, G. (1981). Pragmatics and conversational rhetoric. In H. Parret, M. Sbisà, and J. Verschueren (Eds.), *Possibilities and limitations of pragmatics*, pp. 413–442. Amsterdam: John Benjamins.
- Levinson, S. C. (1983). *Pragmatics*. Cambridge textbooks in linguistics. Cambridge University Press.
- Lewis, D. (1988). Relevant implication. *Theoria* 54(3), 161–174.
- Matsumoto, Y. (1995). The conversational conditions on horn scales. *Linguistics and Philosophy* 18, 21–60.
- McCawley, J. (1978). Conversational implicature and the lexicon. In P. Cole (Ed.), *Pragmatics*, Number 9 in Syntax and Semantics, pp. 245–259. Academic Press.
- Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy* 15(5), 509–559.
- Potts, C. (2005). *The logic of conventional implicatures*. Oxford University Press.
- Roberts, C. (1996). Information structure in discourse. In J. Yoon and A. Kathol (Eds.), *OSU Working Papers in Linguistics*, Volume 49, pp. 91–136. Ohio State University.
- Rooth, M. (1985). *Association with Focus*. Ph. D. thesis, University of Massachusetts, Amherst.
- Russell, B. (2006). Against grammatical computation of scalar implicatures. *Journal of Semantics* 23, 361–382.
- Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy* 27(3), 367–391.
- Schulz, K. and R. van Rooij (2006). Pragmatic meaning and non-monotonic reasoning: The case of exhaustive interpretation. *Linguistics and Philosophy* 29(2), 205–250.
- Simons, M., J. Tonhauser, D. Beaver, and C. Roberts (2010). What projects and why. In B. Jackson and T. Matthews (Eds.), *Semantics and Linguistic Theory (SALT) 20*, pp. 309–327.
- Soames, S. (1982). How presuppositions are inherited: a solution to the projection problem. *Linguistic Inquiry* 13(3), 483–545.
- Spector, B. (2007). Scalar implicatures: Exhaustivity and gricean reasoning. In M. Aloni, A. Butler, and P. Dekker (Eds.), *Questions in Dynamic Semantics*, pp. 225–250. Elsevier.
- Wagner, M. (2012). Contrastive topics decomposed. *Semantics and Pragmatics* 5, 1–54.
- Westera (2017). An attention-based explanation for some exhaustivity operators. In *Proceedings of Sinn und Bedeutung*. University of Edinburgh.
- Westera, M. (2013). ‘Attention, I’m violating a maxim!’ A unifying account of the final rise. In R. Fernández and A. Isard (Eds.), *Proceedings of the Seventeenth Workshop on the Semantics and Pragmatics of Dialogue (SemDial 17)*.
- Westera, M. (2014). Exhaustivity through the Maxim of Relation. In Y. I. Nakano, K. Satoh, and D. Bekki (Eds.), *Selected Papers from the Tenth Logic and Engineering of Natural Language Semantics (LENLS 10)*, New Frontiers in Artificial Intelligence, pp. 193–205. Springer.