

Utterances with themes as strategies to address a broad Question Under Discussion¹

Laia MAYOL — *Universitat Pompeu Fabra*

Enric VALLDUVÍ — *Universitat Pompeu Fabra*

Abstract. In a QUD-model of discourse, any utterance elaborates on the maximal QUD in that context. QUDs play an essential role in defining the two parts in which an utterance can be divided: theme and rheme. An utterance must always contain a rheme, which is the part that elaborates on the QUD, and may contain a theme, which replicates content already present in the QUD. Since themes are replicating material already present in the QUD, one may wonder why themes are uttered at all. Vallduví (2016) proposes that themes signal the the QUD update will have an intermediate step and that the QUD being addressed is not the maximal one. In other words, themes mark that the QUD update is non-default. The goal of this paper is to empirically examine one of these non-default updates and, in particular, whether theme-containing utterances can be used to signal that the QUD being addressed is broader than the maximal one and, if so, whether they are necessary in this situation. Two discourse-completion studies in Catalan were carried out. The results show that theme-containing utterances are mostly used to address a broad QUD (as opposed to narrower ones) and that when speakers decide to address a broad QUD, the proportion of theme-containing utterances increases significantly. The use of themes is, however, not required to signal this change of QUD; themeless-utterances can also be used in this context.

Keywords: Question Under Discussion, QUD, theme, rheme, context

1. Introduction

Krifka and Musan (2012) define *information structure* as concerning ‘those aspects of natural language that help speakers to take into consideration the addressee’s current information state, and hence to facilitate the current flow of information’ (2012: p.1), where the addressee’s current information state is taken to function as context of utterance in the relevant sense. Thus, from this perspective, the information structure of a particular utterance U is taken to reflect the manner in which U effects the update of the input context C_1 , the context in which it is uttered.

Central to information structure are the well known utterance-level notions of *theme* and *rheme*, where themes are taken to perform a backward-looking role, somehow anchoring the utterance to C_1 , while the rheme is seen as performing a forward-looking role, acting as exponent of the progression from the input context C_1 to the output context C_2 , which is obtained after U is computed. In other words, the rheme encodes the update potential of U.

The point of any given utterance is to make the conversation progress, i.e., to provide an update for C_1 . In view of this, Vallduví (1992), among others, argues that the rheme is, in fact, the only obligatory element in the theme-rheme partition, whereas themes appear when the backward-looking task they carry out is required in particular conversational situations. In other words, the presence of a theme in any given utterance, i.e. theme-containing utterances, would be flagging

¹We would like to thank the audiences at XPRAG2019, SuB24 and ConSOLE28, where previous versions of this work were presented. We also gratefully acknowledge support from the Spanish State Research Agency (AEI) and the European Regional Development Fund (FEDER, UE) by projects PGC2018-096870-B-I00 and FFI2015-67991-P.

some kind of *nondefault* contextual update, whereas themeless utterances would somehow have a default status.

This is indeed the case in conversational query-answer pairs. The query in a query-answer pair evokes a context C_1 and the answer spells out the update to C_1 . Take for instance the following example:

- (1) a. What are we having for dinner?
b. FISH.

The answer to the query here is a fragment utterance. Fragment utterances, i.e. themeless utterances like *fish*, are in fact the default type of direct answer in a query-answer pair like this one. A theme-containing utterance like ‘We are having fish’ would be marked in a query-answer pair like (1) but, in contrast, becomes necessary in an exchange like (2), since, as mentioned, theme-containing utterances flag that the contextual update is somehow non-default.

- (2) a. What are we having for dinner?
b. FISH.
c. Oh, not again! We have fish several days a week...! Can’t we order pizza?
d. We are having FISH.

Theme-containing (2d) is an answer to query (2a). The theme allows (2d) to be a non-local answer to (2a). In other words, (2d) elaborates on a Question Under Discussion that is not the most salient Question Under Discussion at the time of utterance: the role of the theme is to re-raise the Question Under Discussion to maximality.

Non-local elaborations such as the one in (2) are just one type of non-default update that is flagged by theme-containing utterances. The goal of this paper is to examine the availability of another such non-default contextual update: an update in which the maximally prominent Question Under Discussion is replaced by a related, broader one. The paper is structured as follows: Section 2 presents the background assumed in this paper, sections 3 and 4 present two discourse-completion experiments in Catalan and section 5 concludes the paper with some discussion.

2. Background

In a Question Under Discussion (QUD)-model of discourse (Ginzburg, 1994; Roberts, 1996) any utterance U uttered in a context C_1 elaborates on the maximal QUD in that context (max-QUD_{C_1}). Following Vallduví (2016), QUDs play an essential role in defining the two parts in which an utterance can be divided: theme and rheme. An utterance must always contain a rheme, which is the part that elaborates on the QUD, its actual update potential. It may also optionally contain a theme, which replicates content already present in the QUD. For instance, the explicit QUD in (1a) is answered in (1b) by an all-rhematic fragment utterance. This query could also be answered by an utterance containing a theme, as ‘We are having’ in (2d), which does not address the QUD but merely replicates material from the QUD.

In addition to elaborating on the maximal QUD, an utterance also introduces a new QUD, which is the one that is maximally prominent in the output context. For instance, utterances (1b) or (2d) introduce the QUD ‘Are we having fish for dinner?’, which another participant in the conversation can elaborate on, as shown by the utterance in (3c).

- (3) A: What are we having for dinner?
 B: Fish // We are having fish for dinner.
 C: No, that's not true!

Since themes in theme-containing utterances are replicating material already present in the QUD, it is only natural to wonder why themes are uttered at all. Vallduví (2016) proposes that themes signal the QUD update will have an intermediate step; the QUD the utterance addresses is not the maximal QUD, but a related one. That is, theme-containing utterances promote one particular QUD so that it replaces the one that was maximally prominent in the context. One case in which this occurs is when the maximal QUD is split into several subQUDs, as extensively studied by Büring (2003), among others, and illustrated in (4). Speaker B decides not to address directly the maximal QUD (the explicit query posed by A), but instead divides the QUD into more specific QUDs: namely 'How was the sound?', 'How was the audience?' and 'How was the band?'. In order to identify which QUD is being addressed, the presence of themes ('the sound', 'the audience' and 'the band' respectively) becomes essential.

- (4) A: How was the concert?
 B: *The sound* was awful, but
 the audience was enthusiastic,
 and *the band* was fantastic.

Now, it should follow that the opposite direction is also possible. Themes should also be able to indicate that the utterance is elaborating on a QUD that is broader than the maximal QUD. That is, if a particular issue i_1 is being discussed, then broader issues i_2 and i_3 which include i_1 are also discussed. This is illustrated in (5) (based on a similar example by Steedman (2014)):

- (5) A: Will Anna marry Manny?
 B: Anna HATES Manny.

The maximal QUD introduced by speaker A is whether Anna will marry Manny. The theme-containing utterance produced by speaker B does not directly answer the maximal QUD, but instead introduces a new QUD, namely how Anna feels about Manny. Clearly both QUDs are related and the issue of the relationship between Eva and Manny is broader than the issue of Eva marrying Manny. If we are discussing whether Anna will marry Manny, we are also discussing how she feels about him, and by elaborating on the broader QUD we can also reasonably infer the answer to the explicit query.

While the existence of query-answer pairs such as the one illustrated in (5) is predicted to be possible and discussed by Vallduví (2016) and Steedman (2014), there are no empirical studies we are aware of that have demonstrated that it is indeed possible to answer a query by elaborating on a broader QUD. This paper aims to fill this gap by answering the following two questions: Can themes also signal that the QUD being addressed is a broader QUD than the maximal one? If they can, are they necessary to do so?

We aim to answer these questions by means of two discourse-completion studies in Catalan, where the mapping between syntax and information structural categories is fairly transparent. In particular, themes are encoded at the beginning or end of the clause (Vallduví, 1992). We will consider three main cases in which themes are expressed: (i) when the utterance contains a preverbal overt subject, as in (6a), as opposed to a null subject, as in (6b), (ii) when the ut-

terance contains a preverbal dative occurring with a psychological verb, as in (7a), as opposed to utterances without an overt dative argument, as in (7b), and (iii) when the utterance contains a right-dislocated phrase, as in (8a), or a left-dislocated phrase, as in (8b), as opposed to utterances without dislocation, as in (8c) (see Vallduví (1992) and Villalba and Mayol (2013) for more discussion of the information structure of dislocations in Catalan)².

- (6) a. *El Pep* no beu vi.
the Pep not drink-3SG wine
'Pep does not drink wine.'
- b. No beu vi.
not drink-3SG wine
'He does not drink wine.'
- (7) a. *Al Pep* no li agrada el vi.
to the Pep not DAT like-3SG the wine
'Pep does not like wine.'
- b. No li agrada el vi.
not DAT like-3SG the wine
'He does not like wine.'
- (8) a. No en beu, *de vi*.
not PART_i drinks, of wine_i
- b. *De vi*, no en beu.
Of wine_i not PART_i drinks
'He does not drink wine.'
- c. No en beu.
not PART_i drinks.
'He does not drink it.'

3. Experiment 1

Experiment 1 is a discourse-completion study, in which we manipulated the amount of information available to participants to answer a query. This allowed us to study if there is a correlation between using a theme and addressing a broader QUD.

3.1. Methods

Materials

Participants were asked to answer a query by completing its answer. All answers that participants had to complete started either with 'Yes' or 'No'. In advance, they were given some contextual information and were told that they should answer the queries on the basis of this information.

The experiment has one factor with two conditions. The contextual information participants received to answer the query could be of two types: (i) in the Small Context Condition, participants received only the information necessary to answer the query, while in the Large Context Condition, participants received more information than what was necessary to answer the

²The following abbreviations are used in the glosses: COND (conditional), PART (partitive clitic), SG (singular), SUBJ (subjunctive).

query. In other words, the sentence in the Large Context Condition asymmetrically entails the sentence in the Small Context Condition. (9) illustrates the two conditions and (10) the query and the answer to be completed.

- (9) a. Al Pere no li agrada gens el lluç.
 To the Pere not DAT like-3SG at all the hake
 ‘Pere does not like hake at all.’ [Small Context Condition]
- b. Al Pere no li agrada gens el peix.
 To the Pere not DAT like-3SG at all the fish
 ‘Pere does not like fish at all.’ [Large Context Condition]
- (10) a. Li faig lluç per sopar, al Pere?
 DAT make-1SG hake for dinner, to the Pere
 Should I cook Pere some hake for dinner? [Query]
- b. Sí, ...
 Yes, ... [Answer]

As can be seen, the information in both conditions is sufficient to answer the query. However, in both cases, it takes an inference to do so: if Peter does not like hake, it can reasonably be concluded that I should not cook some hake for Pere, but this is not a logical necessity (i.e. it may be that I should cook Pere some hake, even if he does not like it, because the doctor said that he should eat fish once a day). We choose to pose a query that could be answered by abductive reasoning, as opposed to deductive reasoning (i.e. ‘Does Pere like fish?’), to make the task more natural and interesting. Otherwise, the task may have been too monotonous, which could have affected the productions of our participants.

It should also be noted that all the answers start with a response particle; half of the answers start with ‘Yes’ and the other half with ‘No’. These response particles are rhematic, since they directly address the explicit polar question posed. Response particles are, however, fully compatible with propositional answers, which may vary in their complexity and information structure, i.e. they may contain only rhematic information, just like the answer particle (see (11a)) or incorporate thematic material (see (11b)).

- (11) a. Sí, fes -li.
 Yes, make-2SG DAT
 ‘Yes, you should.’
- b. Sí, al Pere li agrada el lluç.
 Yes, to the Pere DAT like-3SG the hake.
 ‘Yes, Peter likes hake.’

The experiment included 16 critical items, all with the structure illustrated above. Two lists were built, so that each participant only saw each item in one of the conditions. Each list also contained 16 fillers, in which the context contained information from which the answer to the query could be inferred. In those cases, the answer did not start with ‘Yes/No’, but was completely free. The 32 items were presented in randomized order to the participants.

Procedure

The data was collected on an online survey site (<https://spellout.net/ibexfarm>). First, the participants read the consent form and the instructions in which the procedure was explained.

The instructions stated that they would read a sentence and then would be asked to answer a query on the basis of what they had read. They were asked to answer with a full sentence. An example was presented, with several possible answers, some with a theme and some without. Then, each item was presented one at a time. Finally, some data about the linguistic profile of the participants was recorded.

Participants

Thirty-five participants took part in the experiment. They were all native speakers of Catalan and students at the Universitat Pompeu Fabra. As compensation for their participation, they were entered in a raffle to win a gift certificate.

3.2. Data coding

A total of 560 answers (16 items * 35 participants) were collected. One of the authors of the paper coded the answers as either themeless or theme-containing. As explained in section 2, an answer was coded as having a theme if it contained a preverbal overt subject, (12a), a preverbal dative with a psychological verb, (12b), or phrases dislocated to the left or to the right, (12c). If the answer had none of these elements, it was coded as themeless, as the examples in (13).

- (12) a. Ell odia el peix.
he hate-3SG the fish
'He hates fish.'
- b. Al Pere no li agrada gens el peix.
to the Pere not DAT likes at all the fish
'Pere does not like fish.'
- c. No li_i facis lluç, al Pere_j.
not DAT make-SUBJ-2SG hake, to the Pere.
'Do not cook hake for Peter.'
- (13) a. No li agrada.
not DAT like-3SG
'He doesn't like it.'
- b. Detesta el pescado.
hate-3SG the fish
'He hates fish.'

In addition, in the case of data from the Large Context Condition, it was also coded whether the answer addressed the broad QUD ('Does Peter like fish?') or a narrower QUD. We count answers to both the explicit query ('Should I cook Peter some hake for dinner?') and to the query 'Does Peter like hake?' as addressing a QUD narrower than the broad QUD. While examples (12a), (12b) and (13b) would be coded as addressing the broad QUD, examples (12c), and (13a) would be coded as addressing a narrow QUD. One answer has incoherent and, thus, was excluded from the analysis.

3.3. Hypotheses

Based on the previous literature (Vallduví, 2016; Steedman, 2014), we formulate the following hypotheses:

1. More themes are expected in the Large Context Condition than in the Small Context Condition, given that in the latter speakers do not have information to address the broad QUD. In contrast, in the former, they can choose between addressing the broad QUD or some narrower QUDs. We expect that the possibility of addressing the broad QUD will correlate with an increase of the proportion of themes.
2. Within the Large Context Condition, more themes are expected if speakers address the broad QUD than if they address a narrower QUD.
3. Within the Large Context Condition, more answers will address the broad QUD if speakers use a theme than if they do not.

3.4. Results

The results are shown in Figures 1, 2 and 3, each addressing the corresponding hypothesis above. To test for the statistical significance of the results, mixed-effect logistic regressions were performed, using R (R Core Team, 2013) and *lme4* (Bates et al., 2015). All models contained items and participants as random effects. Likelihood ratio tests are used to compare mixed-effects models differing only in the presence or absence of the fixed effect in question.

Figure 1 shows the proportion of themeless and theme-containing utterances in both conditions. As can be seen, there is no difference between the two conditions and around 30% of the utterances present a theme regardless of their conditions. To test for a main effect of condition, a likelihood-ratio test was conducted between mixed-effects models different only in the presence or absence of a fixed main effect of condition. In both models, theme presence/absence was the dependent variable. The likelihood-ratio test showed no difference between the models ($\chi^2 = 1.06, p = .30$).




Figure 1: Proportion of themes in the two conditions

Thus hypothesis 1 is not borne out. However, so far, we have not examined which QUD speakers are really addressing in the Large Context. In order to do so, let us concentrate on the results obtained in the Large Context Condition. Figure 2 shows the proportion of themeless and theme-containing utterances according to which QUD the speaker is addressing within the Large Context. As can be seen, there is a difference between those cases in which the answer addresses the broad QUD and those in which it addresses a narrower QUD: the proportion of themes is larger when the QUD being addressed is the broadest one. To test for a main effect of condition, a likelihood-ratio test was conducted between mixed-effects models differing only in the presence or absence of a fixed main effect of QUD. In both models, theme presence/absence

was the dependent variable. The likelihood-ratio test showed a main effect of QUD ($\chi^2 = 14.53$, $p < .001$). Thus hypothesis 2 is borne out: the presence of themes is significantly higher when speakers address the broad QUD as opposed to a narrower QUD.

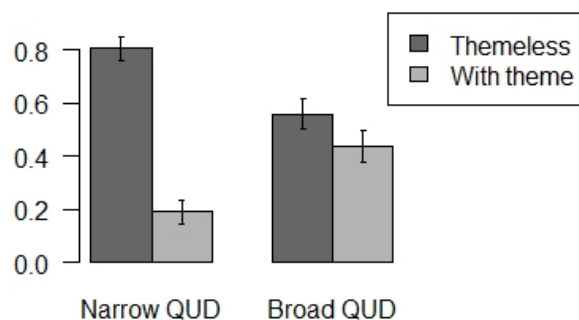


Figure 2: Proportion of themes within the Large Context

Let us now observe the data from the opposite perspective. Figure 3 shows the proportion of broad and narrower QUDs addressed according to the complexity of the utterance (with or without theme). As can be seen, the proportion of broad QUDs increases when the utterance contains a theme. To test for a main effect of condition, a likelihood-ratio test was conducted between mixed-effects models differing only in the presence or absence of a fixed main effect of presence of theme. In both models, type of QUD (broad vs. narrow) was the dependent variable. The likelihood-ratio test showed a main effect of presence of theme ($\chi^2 = 8.14$, $p < .001$). Thus hypothesis 3 is borne out: answers addressing the broad QUD are significantly higher when they contain a theme.

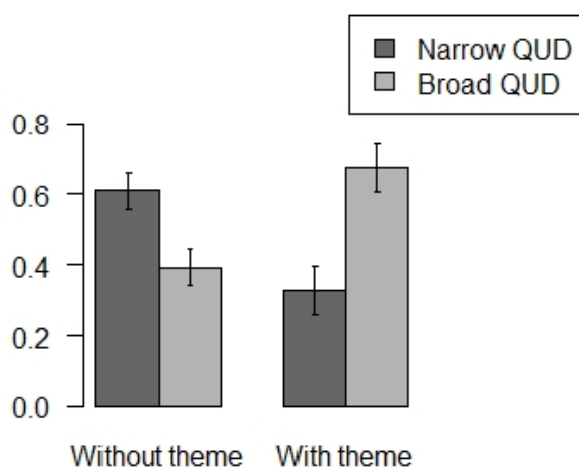


Figure 3: Proportion of QUD type within the Large Context

To sum up the findings in this experiment, (i) when broad QUDs are addressed, the presence of themes increases and (ii) when themes are used, more broad QUDs are addressed. We can also observe that in this latter case (when a theme is used), the majority of QUDs are broad. In contrast, when broad QUDs are addressed, even if the proportion of answer with themes increase, most answers are themeless. That is, themes do not seem necessary to indicate a shift to a broad QUD. We postpone a more in-depth discussion of this result to section 5.

4. Experiment 2

Experiment 2 is another discourse-completion study, which intends to corroborate the findings in Experiment 1 through a different manipulation.

4.1. Methods

Materials

The set-up of the experiment was similar to Experiment 1: participants had to answer a query by completing a sentence on the basis of some contextual information they had previously received. Experiment 2 is different from Experiment 1 in two respects. First, the information participants received was always broader than what was necessary to answer the query (that is, the context was always ‘Large’). Second, there were two types of answer to be completed: while all answers started with either ‘Yes’ or ‘No’, some of the answers also contained a theme. That is, the experiment has one factor with two conditions: Simple Answer, without a theme, or Complex Answer, with a theme. Example (14) shows the context and the query to answer and (15) shows the two types of answer to complete.

- (14) a. La Noemí col·lecciona joies antigues.
the Noemí collect-3SG jewelery antique
‘Noemí collects antique jewelery.’ [Context]
- b. Aquest collaret antic seria un bon regal, per la Noemí?
This necklace antique be-COND-3SG a good gift, for the Noemí
‘Would this antique necklace be a good gift for Noemí?’ [Query]
- (15) a. Sí, ...
‘Yes, ...’ [Simple Answer]
- b. Sí, la Noemí ...
‘Yes, Noemí ...’ [Complex Answer]

The contexts and the queries were the same as used in Experiment 1. Experiment 2 included 16 critical items, along with 16 fillers. Two lists were built, so that each participant only saw each item in one of the conditions. The 32 items were presented in randomized order to the participants.

Procedure and participants

The procedure was the same as explained for Experiment 1. Thirty-three participants took part in the experiment. They were all native speakers of Catalan and students at the Universitat Pompeu Fabra and none of them had participated in Experiment 1. They were entered in a raffle to win a gift certificate.

4.2. Data coding

A total of 528 answers (16 items * 33 participants) were collected. One of the authors of the paper coded whether the answer addressed the broad QUD (‘Does Noemí collect antique jewelery?’) or a narrower QUD. We count answers to both the explicit query (‘Would this antique necklace be a good gift for Noemí?’) and to the query ‘Does Noemí collect antique necklaces?’ as addressing a QUD narrower than the broad QUD. In addition, in the case of

data from the Simple Answer Condition, the answer were coded as either themeless or theme-containing, following the same criteria explained for Experiment 1.

Three answers were excluded from the analysis because they were either incoherent or did not address any of the relevant QUDs.

(16) show some of the answer we obtained together with the coding they received.

- | | | | |
|------|----|---|---------------------------------|
| (16) | a. | Ella col·lecciona joies antigues.
She collect-3SG antique jewellery
'She collects antique jewellery.' | [Broad QUD + theme-containing] |
| | b. | A la Noemí, li agradaria.
to the Noemí DAT like-3SG
'Noemí would like it.' | [Narrow QUD + theme-containing] |
| | c. | Col·lecciona joies antigues.
collect-3SG antique jewellery
'She collects antique jewellery.' | [Broad QUD + themeless] |
| | d. | Seria un bon regal
be-COND-3SG a good gift
'It would be a good gift.' | [Narrow QUD + themeless] |

4.3. Hypotheses

In a parallel fashion to what we did for Experiment 1, we formulate the following hypotheses:

4. We expect that the proportion of answers addressing the Broad QUD will be larger in the Complex Answer Condition than in the Simple Answer Condition. Since complex answers always contain a theme, we expect that this will correlate with more answers addressing the Broad QUD.
5. Within the Simple Answer Condition, more answers will address the broad QUD if speakers uses a theme than if they do not.
6. Within the Simple Answer Condition, more themes are expected if speakers address the broad QUD than if they address a narrower QUD.

4.4. Results

The results are shown in Figures 4, 5 and 6, each addressing the corresponding hypothesis above. The same method described for Experiment 1 was used to test for the statistical significance of the results.

Figure 4 shows the proportion of broad and narrow QUDs in both conditions. In the Complex Answer Condition, the proportion of answers addressing the broad QUD is slightly higher (70%) than in the Simple Answer Condition (61%). To test for a main effect of condition, a likelihood-ratio test was conducted between mixed-effects models differing only in the presence or absence of a fixed main effect of condition. In both models, the QUD addressed (broad vs. narrow) was the dependent variable. The likelihood-ratio test showed a main effect of the condition ($\chi^2 = 7.06$, $p < .01$). Thus hypothesis 4 is borne out.

Let us now concentrate in the results of the Simple Answer Condition so that we can examine




Figure 4: Proportion of QUD type in the two conditions

the types of answers speakers produced and whether they contained a theme or not. Figure 5 shows the proportion of broad and narrower QUDs addressed according to the complexity of the utterance (with or without theme). As can be seen, the proportion of broad QUDs increases when the utterance contains a theme. To test for a main effect of condition, a likelihood-ratio test was conducted between mixed-effects models differing only in the presence or absence of a fixed main effect of presence of theme. In both models, type of QUD (broad vs. narrow) was the dependent variable. The likelihood-ratio test showed a main effect of presence of theme ($\chi^2 = 8.39$, $p < .01$). Thus hypothesis 5 is borne out: the proportion of answers addressing the broad QUD is significantly higher when they contain a theme.

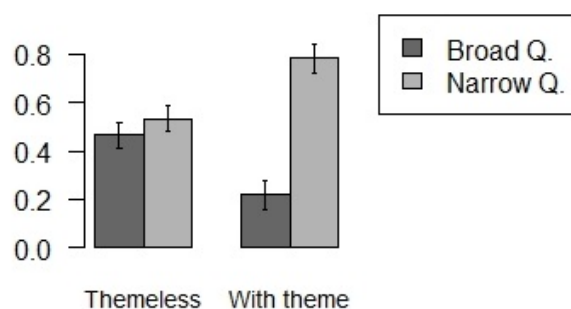


Figure 5: Proportion of QUD type within the Large Context

Finally let us now look at the data from the opposite perspective. Figure 6 shows the proportion of themeless and theme-containing utterances according to which QUD the speaker is addressing within the Simple Answer Condition. As can be seen, there is a difference between those cases in which the answer addresses the broad QUD and those in which it addresses a narrower QUD: the proportion of themes is larger when the QUD being addressed is the broadest one. To test for a main effect of condition, a likelihood-ratio test was conducted between mixed-effects models differing only in the presence or absence of a fixed main effect of QUD. In both models, theme presence/absence was the dependent variable. The likelihood-ratio test showed a main effect of QUD ($\chi^2 = 10.97$, $p < .001$). Thus hypothesis 6 is borne out: the presence of themes is significantly higher when speakers address the broad QUD as opposed to a narrower QUD.

The findings in Experiment 2 largely replicate those of Experiment 1: (i) when broad QUDs are addressed, the presence of themes increases and (ii) when themes are used, more broad QUDs are addressed. As found in Experiment 1, we can also observe that themes do not seem to be




Figure 6: Proportion of themes within the Large Context

necessary to address a broad QUD. That is, most of the answers addressing a broad QUD are themeless, as can be seen in Figure 6.

5. Discussion and conclusion

In this paper, the role of theme-containing utterances has been examined in connection to their ability to signal that the QUD being addressed is broader than the maximally prominent QUD at a particular point of the discourse. Two discourse-completion studies were carried out which confirm that, indeed, there is a relationship between presence of a theme and elaboration of a broad QUD. Participants produced more themes when addressing a broad QUD, as opposed to a narrower QUD. In addition, they also mostly addressed the broad QUD when their answer contained a theme.

We can, thus, conclude theme-containing utterances can flag the QUD being addressed is broader than the maximal one. However, it is also obvious that there is significant variation in the data and that the correlation between themes and broad QUDs is far from being a categorical one. One very striking result in both experiments is that, when broad QUDs are addressed, most answers are themeless (even if the proportion of answers with themes increases significantly). That is, while themes are a mechanism that can signal a shift to a broad QUD, they are not necessary to do so: it is possible to change the QUD even in absence of a linguistic marker. We believe that themes will tend to be absent when the broad QUD is deemed to be highly predictable by the speaker. That is, when a speaker is asking whether she should cook hake for Peter, it is obvious that it may be relevant to discuss Peter's tastes. Therefore, the QUD 'Does Peter like fish?' is relevant for the conversation and, in effect, predictable. Given that the role of the theme is to prepare the context for an unexpected QUD, the more predictable the QUD is (even if it is not the maximally prominent one), the less necessary the presence of themes should be. Thus, the high number of themeless utterances in our data should not be surprising, given that in many occasions the QUDs they address can be considered to be highly predictable.

Related to this last point, it is worth mentioning that, looking at the data item by item, we can observe how the proportion of answers addressing the broad QUD varies greatly. Averaging the results from the two experiments, the proportion of answers addressing the broad QUD ranges from 16.5% to 87% depending on the item. That is, we find items on both ends of the scale (items for which the QUD addressed is almost always the broad one and items for which the QUD addressed is almost always the narrow one). Furthermore, the item distribution is fairly balanced throughout the scale, as can be observed in Table 1.

	# of items
Less than 25%	3
Between 25% and 50%	6
Between 50% and 75%	4
More than 75%	3

Table 1: Number of items according to the % of answers addressing the broad QUD.

To understand the source of this distribution, we can observe the three items that triggered more answers to the broad QUD in (17) and the three items that triggered less answers to the broad QUD in (18). In both cases, the query is preceded by the Large Context in parentheses.

- (17) a. [Pere does not like fish.]
 Li faig lluç, al Pere, per sopar?
 ‘Should I cook Pere some hake for dinner?’
- b. [Carles never drinks alcohol.]
 Li porto una cervesa, al Carles?
 ‘Should I bring Carles a beer?’
- c. [Andreu hates touristy beach towns.]
 Quedem amb l’Andreu a Salou?
 ‘Should we meet Andreu in Salou?’
- (18) a. [Roser used to be allergic to berries, but not anymore].
 Li dono maduixes per postres, a la Roser?
 ‘Should I give Roser strawberries for dessert?’
- b. [Sònia has been at the Louvre Museum multiple times.]
 Li proposo d’anar a París, a la Sònia?
 ‘Should I ask Sònia if she wants to go to Paris?’
- c. [Ariadna is at school every evening from 6 to 8.]
 ‘Li dic a l’Ariadna si vol anar al cine el dijous a les 6:30?’
 Should I ask Ariadna if she wants to go to the movies on Thursday at 6.30?

In all items, the propositional content of the context entails a proposition by which we can infer the answer to the explicit query (i.e. ‘Carles never drinks alcohol’ entails ‘Carles never drinks beer’, which is useful to answer the query in (17b)). Then, why do we find such a different rate of answers addressing the broad QUD in (17) and (18)? We believe that the difference lies in the problem that (our participants imagined that) the questioner is trying to solve when she poses a query. For instance, when the questioner poses a query such as (17a), in fact, the problem she is trying to solve is what to give Pere for dinner. Under these circumstances, the person who is answering the query (in the Large Context) not only knows that Pere does not like hake, but has additional information that might be useful to solve the problem of the speaker: Pere does not like fish in general and this information is useful to eliminate other relevant alternatives. By offering an answer which elaborates on the broad QUD, the answerer is helping the questioner to find a solution to her problem. Something similar happens in (17b) and (17c), where the problems that we can imagine the questioner is trying to solve are what drink should she offer to Carles and where they can meet with Andreu, respectively. By giving an answer that addresses the broad QUD, the answerer is eliminating relevant alternatives and

giving an answer that is more informative and helpful to resolve the problem of the questioner. In contrast, in the items of (18), either it is not as obvious which problem the questioner is trying to solve or the contextual information is not relevant to solve it. In (18a), given that the answer is positive (that is, Roser can eat strawberries), the problem is solved just by addressing the narrow QUD and it is not relevant that the answerer gives more information than necessary by addressing the broad QUD. In (18b), the information that Sònia has been in the Louvre Museum may eliminate Paris as a destination, but is not helpful to eliminate other alternatives. Finally, in (18c), it is possible that participants imagined a scenario in which the interlocutors have already decided to go to the movies at a particular day and time and they just need to find out whether Ariadna can join them or not. In this scenario, the details about Ariadna's schedule are not relevant or helpful to solve the problem. To sum up, which QUD participants in a conversation decide to address at each moment crucially depends on the goal they believe the questioner has in posing the query.

To sum up, in this paper, we experimentally explored the understudied connection between the presence of a theme in an utterance and its ability to address a QUD broader than the most prominent one. The results for Catalan show that this connection clearly exists. One may wonder to what extent our results can be extended to other languages and, in particular, to languages without null pronouns, such as English. For a language like English, we would predict that we should find a difference in the type of subject present in the utterance: utterances with unstressed pronouns should behave like null pronouns and, thus, be more likely to answer a narrow QUD, while the presence of a more complex referential expression should correlate with an increase of elaborations of the broad QUDs. We leave it for future work to replicate this type of studies in other languages.

References

- Bates, D., M. Mächler, B. Bolker, and S. Walker (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67(1), 1–48.
- Büring, D. (2003). On D-Trees, Beans, and B-Accents. *Linguistics & Philosophy* 26(5), 511–545.
- Ginzburg, J. (1994). An update semantics for dialogue. In *Proceedings of the first International Workshop on Computational Semantics*, Tilburg, The Netherlands.
- Krifka, M. and R. Musan (2012). Information structure: Overview and linguistic issues. In M. Krifka and R. Musan (Eds.), *The expression of information structure*, pp. 1–44. Berlin: De Gruyter.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Roberts, C. (1996). Information structure in discourse: Towards an integrated formal theory of pragmatics. *OSUWPL Volume 49: Papers in Semantics*, 91–136.
- Steedman, M. (2014). The surface-compositional semantics of English intonation. *Language* 90(1), 2–57.
- Vallduví, E. (1992). *The Informational Component*. New York: Garland Publishing.
- Vallduví, E. (2016). Information structure. In *Cambridge Handbook of Formal Semantics*. Cambridge: Cambridge University Press.
- Villalba, X. and L. Mayol (2013). Right-dislocation in catalan: Tails, polarity and activation. *International Review of Pragmatics* 5(1), 87–117.