Shared mechanism underlying unembedded and embedded enrichments: Evidence from enrichment priming

Chao SUN — *University College London*
Richard BREHENY — *University College London*

Abstract. In this paper, we use a priming paradigm to explore the mechanisms underlying unembedded and embedded scalar enrichments. In particular, the aim is to see if local pragmatic enrichment could be a shared mechanism, involved in both. The two experiments presented adopt Bott & Chemla's (2016) enrichment priming paradigm and test whether unembedded and embedded enrichments could prime each other. The goal is to investigate whether local pragmatic enrichment is indeed being accessed for the interpretation of the unembedded scalar and whether local enrichments, like other lexical semantic phenomena, are susceptible to priming.

Keywords: pragmatics, scalar enrichments, priming.

1. Introduction

Scalar implicatures are widely discussed as potentially Gricean conversational implicatures. (1-2) are examples of scalar implicatures, where what follows '□>' are implications that would follow in many easily imaginable situations:

1. Player A scored some of his shots.
   □> Player A did not score all of his shots.
2. A: Alice was planning to cut the grass and wash the car. I wonder how she got on.
   B: She cut the grass.
   □> Alice didn't wash the car.

Here we use 'scalar implicature' and 'scalar enrichment' as descriptive terms for the phenomenon where an implication arises which seems to involve the negation of a contextually salient alternative.

Many well-known proposals explain the implications in (1-2) broadly-speaking as Gricean conversational implicatures (see e.g. Gazdar, 1979; Geurts, 2010; Sauerland, 2004). On this kind of approach, an alternative for the assertion is inferred to be not true on the basis of reasoning about the speaker's intentions. A widely discussed limitation of this approach is that it cannot explain certain so-called 'embedded scalar enrichments' (Chierchia, 2004; Chierchia, Fox, and Spector, 2012; Potts et al., 2016). An example of an embedded enrichment is given in (3) – taken from Potts et al. (2016):

3. Exactly one player hit some of his shots
   □> Exactly one player hit some and not all of his shots

The observation is that it is unclear how a Gricean account of contextual implications can derive this effect, since the implication under (3) neither entails nor is entailed by what might have been literally asserted in (3). It seems rather that this effect is the product of an

Proceedings of Sinn und Bedeutung 22, vol. 2, ZASPiL 61, pp. 425–441. ZAS, Berlin.

enrichment of a sub-constituent of (3) (Chierchia, 2004) and it patterns together with other ‘local pragmatic effects’ that have been discussed virtually since Grice’s original theory was proposed (Cohen, 1971; Wilson, 1975; Carston, 1998).

Perhaps unlike any other kind of ‘local’ or ‘embedded’ pragmatic effect, embedded scalar enrichments have been quite intensively studied experimentally, with the aim largely being to establish the extent to which embedded scalar enrichments actually arise. Results have varied quite widely (Geurts and Pouscoulous, 2009; Chemla and Spector, 2011), and there has been some critical discussion of the methods used and the interpretation of results. However, a recent study reported in Potts et al. (2016) was designed to avoid many of the perceived methodological flaws of previous research, and found that participants quite regularly responded to a task based on an understanding of sentences like (3) as involving the implication indicated above.

1.1 Theoretical background

Two rather different approaches to embedded scalar phenomena have been outlined in the literature. According to the Grammatical Theory of Scalars (GT), the effects described in (1-2) and (3) are explained by the presence of a operator in the syntactic representation for the sentence. The only difference between the embedded scalar enrichment in (3) and the unembedded enrichments in (1-2) is the scope site at which the operator is inserted. This difference is illustrated in (4-5), where (4) is the LF for (1) and (5) is the LF for (3). Here *O* is an operator whose interpretation relates that of its argument and the argument’s scalar alternatives in a manner akin to ‘only’ (see Chierchia et al., 2011 for details):

4. [O [[Player A]$_i$ [t$_i$ hit some of his shots]]]
5. [[Exactly one player]$_i$ [O [t$_i$ hit some of his shots]]]

According to GT then, we can say that there is but one operation by which both unembedded and embedded scalar enrichments are derived.

Somewhat in contrast to GT, a variety of more-or-less Gricean approaches see at least some embedded scalar enrichments as the result of a separate process of local adjustment to the literal meaning of expressions. This approach takes a cue from the research tradition mentioned above in that it sees embedded scalar enrichments as a result of a general local enrichment mechanism that can result in a variety of different embedded effects, not just scalar enrichments (see Carston, 2002). According to this approach, it is conceptually possible that even unembedded scalar enrichments result from local enrichment. However, it is also allowed that unembedded scalar enrichments could be the result of general reasoning about the speaker’s intentions, along the lines of the well-known Gricean approach to scalars.

A recent articulation of this view is presented within the RSA framework (Frank and Goodman, 2012; Goodman and Frank, 2016). In that framework, it is possible to explain unembedded scalar implicatures in terms of general reasoning speakers and hearers may engage in about each other, making assumptions about how speakers would optimise the utility of their utterance by making the most specific assertion compatible with their knowledge. In addition, as Bergen et al. (2016) observe, it is possible to incorporate the apparent fact that local enrichments of an expression's literal meanings are possible. Bergen et al. set out a framework for computing an interpretation of an utterance given that expressions may be interpreted using their literal meaning or one of a number of possible enrichments. Thus a sentence such as (1), containing an unembedded 'some', may imply *not all* because this can be inferred by 'global' reasoning about the speaker, as set out in the standard RSA approach; alternatively, the implication may simply arise as an entailment of the locally enriched interpretation of 'some'. Bergen et al.'s RSA with lexical uncertainty (RSA-LU) simply builds this fact into the reasoning that speaker and hearer engage in. Likewise, when 'some' appears in an embedded context like (3), the framework simply takes into account that there are several logically independent readings available.

Potts et al. (2016) show that models derived from RSA-LU better predict the results of an experiment in which participants are asked to judge sentences like (3) against visual displays that make the unenriched and locally enriched interpretation true. Potts et al. observe that model performance can be closer or further from actual participants' responses depending on how the prior probabilities of local enrichments are adjusted. This point will be relevant to our discussion of the results of our experiments below. For now, it is sufficient to observe that RSA-LU is a framework for explaining embedded and unembedded scalar implicatures (as in (1-3)) where a single operation (lexical enrichment) is active in both cases, but where there is a second operation (global reasoning) in the unembedded case.

Thus, two approaches suggest that a common means exists for deriving unembedded and embedded scalar enrichments. In this paper, we utilise the 'enrichment priming' paradigm of Bott & Chemla (2016) as a means to determine experimentally whether, in fact, embedded and unembedded scalar enrichments share a mechanism, or have a common operation.

1.2 Enrichment priming paradigm

Bott & Chemla (2016) developed an enrichment priming paradigm for the purpose of obtaining empirical evidence for shared mechanisms within and across different categories of unembedded scalar enrichments (i.e. quantifiers, numerals, ad hoc). In this task, each sentence is presented with two pictures, and participants are asked to click on the picture that is a better match for the given sentence. The critical items for a 'within-category' priming condition are illustrated in Figure 1.

In this condition, the target and prime trials involve the same enrichment category. That is, a target trial with 'some' is preceded by prime trials also with 'some'; a target trial with numerals is preceded by prime trials with numerals, etc. There are two types of prime trials, Strong and Weak. Consider *some → some* in the top panel of Figure 1. In the Strong prime condition, given the sentence *Some of the symbols are clubs*, the 'strong' image shows some and not all symbols are clubs, and the other, 'weak' image shows all symbols are clubs. The strong image makes the scalar-enriched interpretation (*some and not all symbols are clubs*) true. The 'weak' image is only true on an unenriched interpretation of the target sentence. Participants who choose the strong image prior to a target trial are thus primed by the SI-enriched reading. In contrast, in the Weak prime condition, given the sentence *Some of the* *symbols are stars*, one picture contains all stars and the other contains only non-stars. Neither picture makes the interpretation that includes the scalar implicature true. Participants who give a correct response in Weak prime trials have had to entertain the unenriched interpretation of the sentence prior to the target trial.

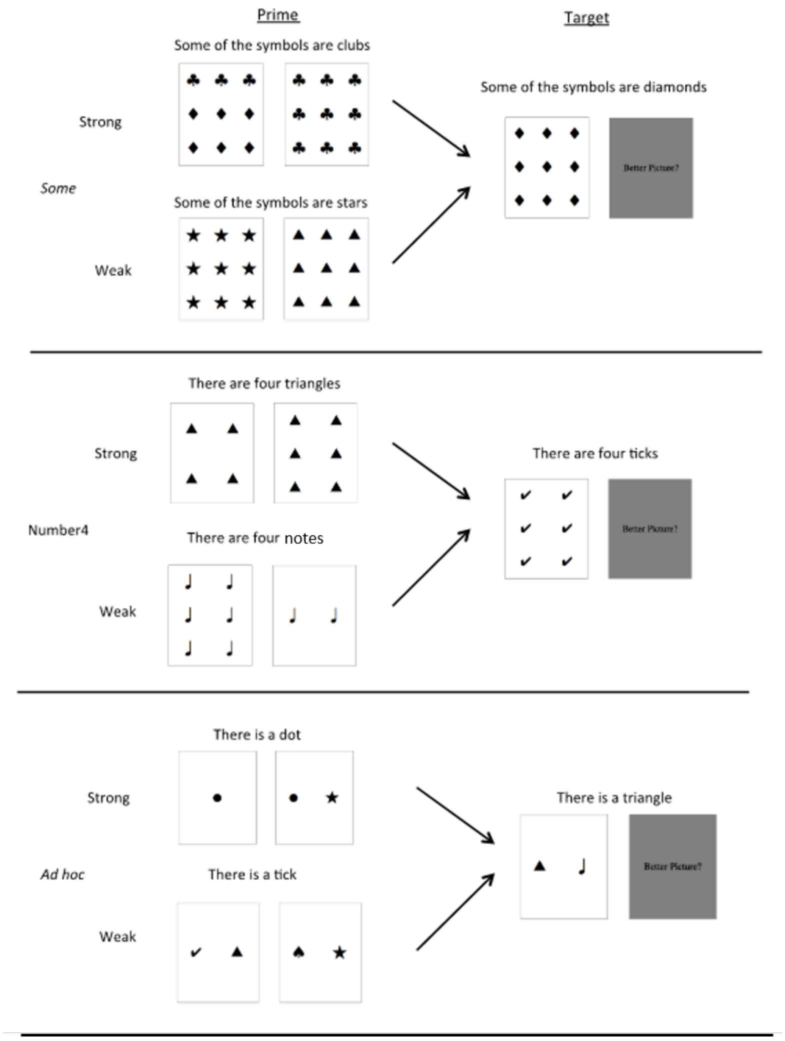

Figure 1 *Example items in Bott & Chemla (2016)*[1]

For the target trials, Bott & Chemla (2016) adopted the 'Better-picture' method used in Huang, Spelke & Snedeker (2013). Participants are shown one of two images while the other is covered. Participants are told that if they think that the covered picture would be a better match for the sentence, they can choose the covered picture. In this design, the visible image makes the unenriched reading true. Since the visible picture is inconsistent with the SI-enriched reading of the target sentence, choosing the covered picture indicates that participants access the SI-enriched reading.

[1] We have had permission from Dr Bott to use the figure from Bott & Chemla (2016).

In addition to within-category priming, the other condition is between-category priming, where the target and prime trials involve different enrichment categories. For instance, a target trial with number term (e.g. 'four') is preceded by prime trials with 'some'. Bott & Chemla included all between-scale combinations in this condition, such as *some* ↔ *number*, *some* ↔ *ad hoc*, and *number* ↔ *ad hoc*.

The logic behind this paradigm is that, if there is a shared derivation mechanism which is subject to priming, then for both conditions it is more likely for participants to access the enriched reading of the target sentence (i.e. choosing the covered picture) after strong prime trials than after weak prime trials. Their results show a within-category priming and a between-category priming effect. The within-category effect was stronger. There was also a surprising effect of within / between, such that more SI-based responses occurred in the between-category condition. We will return to discuss the latter result in Section 4. The main result of this kind of study, however, is that unembedded scalar enrichments can be primed by unembedded scalar enrichments.

Bott & Chemla (2016) interpreted the between-category priming effect as evidence for activation of shared mechanisms in deriving enrichments involving different scales. As for the within-category priming effect, they suggested that along with the activation of the derivation mechanism, there could also be a lexical priming effect, which is an association between the stimulus, the derivation mechanism and specific alternative. For Bott & Chemla, the between-category priming effect is most interesting result, because it shows that general SI derivation mechanism can be primed.

The general idea, then, is that an enrichment priming paradigm could be employed to investigate whether local pragmatic enrichment is a shared mechanism between unembedded and embedded scalar enrichments.

1. Experimental overview

The first goal of the experiments in this paper is to determine whether embedded and unembedded scalar phenomena have a shared mechanism. We investigate the mechanisms underlying unembedded scalar enrichment using the same paradigm as in Bott & Chemla (2016). The rationale is that, if unembedded scalar implicatures are derived using an operation or mechanism that is also involved in embedded scalar enrichments, then participants should be more likely to access embedded enrichments after strong primes with unembedded scalar implicature than after weak primes with no implicature. The critical items are illustrated in Figure 2.

In the embedded target condition, the target trial involving embedded 'some' is preceded by prime trials involving unembedded 'some'. In strong primes, the unembedded scalar implicature is true, while in weak prime trials, the unembedded scalar implicature is false. For example, given a prime sentence 'Some of the symbols are diamonds', in strong primes, the sentence is presented with one picture depicting a row with some but not all symbols being diamonds and another picture depicting a row with all symbols being diamonds. The 'some-not-all' picture makes the SI-enriched reading true. For the weak primes, the same sentence is presented with one picture in which all symbols in the row are diamonds and one picture depicting a row of non-diamond symbols. Neither picture makes the SI-enriched reading true. Thus, participants are primed by the SI-enriched reading in strong primes and the unenriched reading in weak primes.

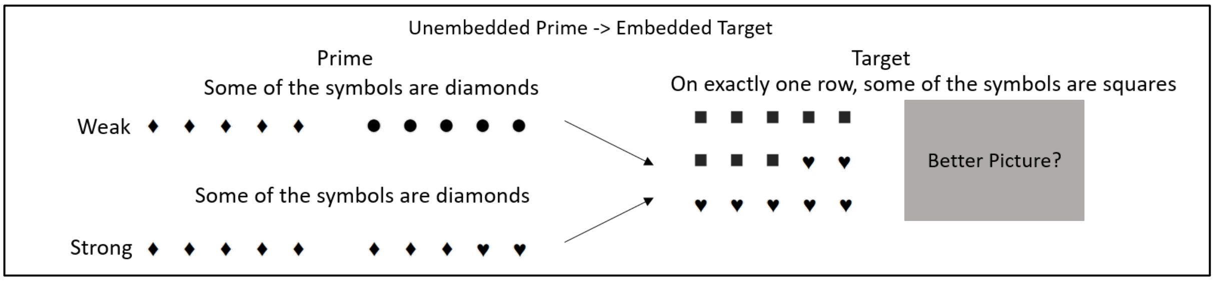

Figure 2 *Critical items for embedded target condition in Experiment 1 and 2*

As in Chemla & Bott, we employ the covered picture paradigm in the target trials. We have experimental trials in which a sentence with an embedded scalar term is the target. We also include a set of trials where an unembedded sentence is the target, following embedded prime trials. For target trials in the embedded target condition, a target sentence like 'On exactly one row, some of the symbols are squares' is presented with a visible picture and a covered picture. The visible picture makes the locally enriched reading true and other available readings false. The image in Figure 2 shows the visible image having two rows containing squares. One of those has some and not all squares, the other has all squares. Only if the sentence is understood as *On exactly one row, some and not all of the symbols are squares* would a participant not choose the covered card. If the literal meaning of the target sentence is accessed, or even an interpretation that includes a global implicature, the participant should choose the covered card.

This is a change from Bott & Chemla's procedure. As previously mentioned, the visible picture used in Bott & Chemla's paradigm makes the literal reading true and SI-enriched reading false. The motivation for changing their design comes from the availability of the global-SI reading. The global-SI reading of the target sentence is that *on exactly one row, some symbols are squares and it's not true that on exactly one row, all symbols are squares*. If the target sentence is presented with a visible picture that makes the literal reading true, as shown in Figure 3 (left), then participants might choose the covered picture because they derive a reading of the sentence that includes a global SI and expected a better match, such as Figure 3 (right). If this is the case, then choosing the covered picture in Figure 3 might reflect a mixture of local reading and global reading.

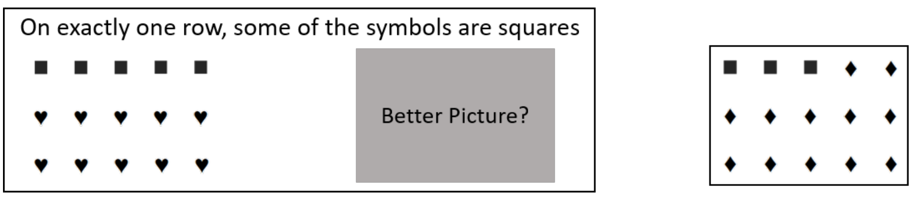

Figure 3 *Alternative displays*. The target (left image) consists of one picture that makes the literal reading true and the 'Better Picture?' option. The right image makes the global-SI reading of 'On exactly one row, some of the symbols are squares' true.

Thus, in order to properly measure the rate of locally enriched reading, in both Experiments 1 and 2 below, the embedded target sentence is paired with a visible picture for which the sentence is false on any available reading except for the local one. In this case, choosing the visible picture indicates that participants access the locally enriched reading, whereas choosing the covered picture indicates that they access either the literal reading or the global reading.

Regarding whether unembedded enrichments could prime embedded enrichments, the grammatical account predicts a priming effect, as there is a single mechanism for both prime and target trials involving *O* operator in LF. On the other hand, the RSA-LU approach predicts priming between the two based on the mechanism of lexical adjustment, which can be used in both prime and target trials. However, RSA-LU does not rule out the possibility that there is no priming effect. This is so since there are two mechanisms underlying scalar enrichments, rather than a single one. It is possible that the lexical adjustment mechanism is not used very much in prime trials. If this is the case, then there might not be a priming effect between unembedded and embedded enrichments.

In addition to the embedded target condition, both experiments also included an unembedded target condition. In the unembedded target condition, the target trial involving unembedded ‘some’ is preceded by prime trials involving embedded ‘some’. Experiment 1 and 2 differ in the prime items used in unembedded target condition, which will be discussed in more detail below. Regarding whether embedded enrichments could prime unembedded scalar implicature, the grammatical account again predicts a priming effect on the basis of a single shared mechanism. The RSA-LU also predicts a priming effect, as the lexical adjustment mechanism is needed for embedded prime trials (especially in Experiment 2), and the target trial can be enriched in the same way.

2. Experiment 1

3.1. Overview and prediction

In prime trials, participants were presented with a sentence paired with two pictures. Their task was to click on the picture that makes the sentence true. The sentences contained a scalar term ‘some’, which could occur in either unembedded or embedded position. Three types of pictures were available for each sentence: (i) false pictures, which make all possible readings false, (ii) weak pictures, which make the literal reading true but the enriched reading false, and (iii) strong pictures, which make enriched readings true. As will become clear below, the design of this study differs a little from Bott & Chemla. In their paper, strong pictures make not only the enriched meaning true but also the literal meaning. This is also the case in our unembedded prime and target trials, as well as the embedded prime trials in Experiment 1. However, it is not the case for the embedded target trials in either Experiment 1 or Experiment 2, for the reason discussed above (in relation to Figures 2 & 3). As mentioned above, in order to avoid responses that were not solely based on a genuine local enrichment operation, we had to make the verifying scenario for the embedded target sentence falsify the literal meaning.

Two types of priming effects were examined: unembedded prime → embedded target, as shown in Figure 2, and embedded prime → unembedded target, as shown in Figure 4 below. There were two types of prime trials. Participants were primed by the literal reading in weak primes and the enriched reading in strong primes. Following the procedure in Bott & Chemla (2016) and Raffray & Pickering (2010), each target trial was preceded by two prime trials, in order for the priming effect to be given a better chance of having an effect. For target trials, the sentence was presented with an open picture and a covered picture. Participants were instructed to click on the covered picture ('Better Picture?') if they thought there was a picture that would be a better match for the given sentence.

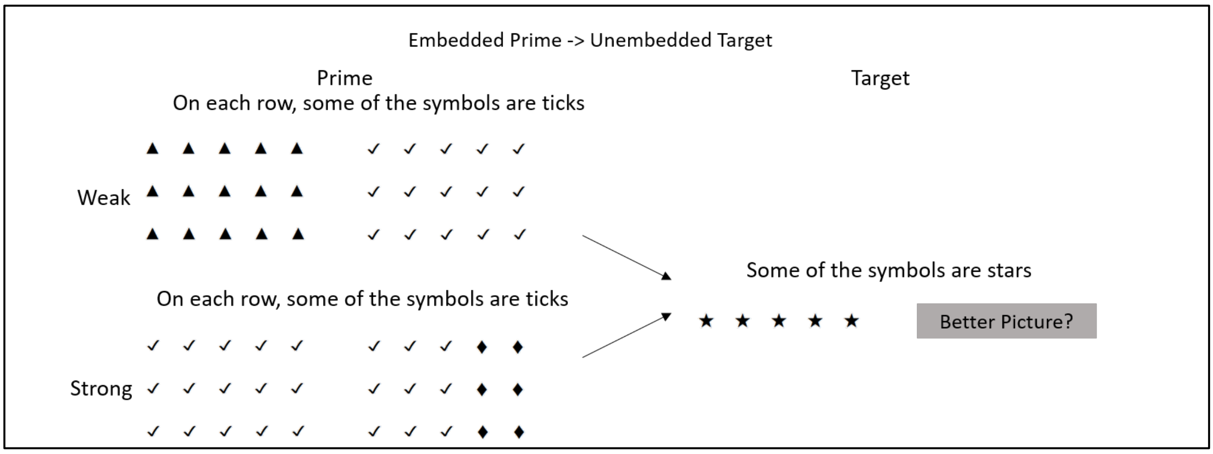

Figure 4 *Critical items for unembedded target condition in Experiment 1*

The embedded target condition has been discussed in detail in the previous section. Here we focus on the unembedded target condition. The critical items of this condition are illustrated in Figure 4. In the unembedded target condition, the target trial involving unembedded 'some' was preceded by prime trials involving embedded 'some'. For embedded prime trials, given the prime sentence like 'On each row, some of the symbols are ticks', in strong primes, the sentence was presented with a weak picture depicting all symbols being ticks and a strong picture depicting rows of symbols with some but not all being ticks. The strong picture made the locally enriched reading of the sentence true (i.e. *On each row, some but not all of the symbols are ticks*). For the same sentence, in weak primes, it was presented with a weak picture and a false picture depicting all symbols being non-ticks. Neither picture made the local reading true. Participants were thus forced to access the literal reading in weak primes.

Note that the sentences used for embedded target trials like 'on exactly one row, some of the symbols are squares' were not used in embedded prime trials. This is because when 'some' is embedded under a non-monotonic quantifier, the literal reading and local enriched reading are logically independent. Thus, if non-monotonic cases are used as embedded primes, there is no better picture (in the sense of entailment) between a picture that makes the literal reading true and a picture that makes the enriched reading true.

As for unembedded target trials, the target sentence was the same as the one used for unembedded prime trials. Unlike embedded target trials, here the unembedded target sentence was presented with a visible picture that made the literal reading true. In this case, choosing the visible picture indicates that participants access the literal reading, whereas choosing the covered picture indicates that they access the SI-enriched reading.

In general, both the GT and the RSA-LU approach predict priming effects between unembedded and embedded enrichments, since both approaches assume there is a shared mechanism between unembedded and embedded enrichments. Overall, the rate of enriched-reading responses to target trials should be higher after strong primes than after weak primes. However, as mentioned above, there is a subtle difference between the two approaches in terms of the potential strength of priming in the different target conditions. The GT says that there is only one mechanism of exhaustification and it is present in both unembedded and embedded scalar enrichments. Thus, whether unembedded trials or embedded trials are primes, the subsequent target should receive more enriched responses after strong prime trials. For the RSA-LU approach, this prediction holds for the embedded prime → unembedded target trials. However, for the case where the prime is unembedded, there are two routes to an enriched response. Only if enriched responses in unembedded primes involve a local pragmatic enrichment should there be substantial priming in the embedded target conditions. We shall return to this difference below.

3.2. Method

3.2.1. Participants

20 participants were recruited via Prolific Academic (http://prolific.ac). All participants were native English speakers.

3.2.2. Materials

This experiment had a two-by-two within-participant design. The two independent variables were the embeddedness of the target and the type of the prime. These two variables generated four prime-target combinations, as shown in Table 1. Sixteen experimental prime-target triplets were constructed. In each triplet, one target trial was preceded by two prime trials. Each trial consisted of a single sentence and two pictures. Eight triplets formed the unembedded prime → embedded target trials, the other eight formed the embedded prime → unembedded target trials. In half of the unembedded prime → embedded target trials, the target was preceded by two weak primes, while in the other half, the target was preceded by two strong primes. This was the same for the embedded prime → unembedded target trials.

Target embeddedness	Prime type	Number of sets	Number of trials
embedded target	weak	4	12
	strong	4	12
unembedded target	weak	4	12
	strong	4	12
			48

Table 1 *Design of experimental items in Experiment 1*

For unembedded prime and unembedded target trials, the sentence was of the form *Some of the symbols are [symbol]*. For embedded prime trials, the prime sentence was of the form *On each row, some of the symbols are [symbol]*, whereas for embedded target trials, the target sentence was of the form *On exactly one row, some of the symbols are [symbol]*. The symbols came from the set of circles, crosses, diamonds, hearts, squares, stars, ticks, and triangles.

48 filler trials were constructed. As with experimental trials, each consisted of a single sentence and two pictures. The sentence either contained 'some' as in *Some of the symbols are [symbol]* or *On each row, some of the symbols are [symbol]*, or contained 'all' as in *All of the symbols are [symbol]* or *On each row, all of the symbols are [symbol]*. Following the design in Bott & Chemla (2016), each type of filler sentences occurred in three situations: (i) the sentence was presented with a strong picture and a 'Better Picture?', (ii) the sentence was presented with a false picture and a 'Better Picture?', and (iii) the sentence was presented with a false picture and a strong picture. (i) and (ii) were included to counterbalance the times when, in the target trials, the covered picture ('Better Picture?') was always paired with the weak picture. These trials also counterbalanced the extra times when in prime trials the sentence was always paired with two visible pictures. (iii) was included so that all possible pair combinations of three picture types (false, weak, strong) occurred equally frequently.

In total, Experiment 1 contained 48 experimental trials (i.e. 16 prime-target triplets) and 48 fillers. The triplets of trials and the fillers were presented in a randomized order created for each participant. For prime trials, the position of the correct choice was counterbalanced across trials, so that for half of the trials the correct choice was on the left, and for the other half on the right[2]. Furthermore, for half the experimental triplets the correct choice was on the same side for the first and the second prime, while for the other half it was on the opposite side. For target trials, the covered picture was always on the right. In addition, in one dual prime-target triplet, a different symbol was used as the predicate for each sentence. There were 8 symbol types. Each was used as the predicate an equal number of times.

3.2.3. Procedure

Participants were instructed to click on the picture that made the sentence true. On occasions where one of the two pictures were covered, the task was the same. But participants were told that "if you think that there is a picture that would be a better match for the sentence, click on the 'Better Picture?' option". Two examples were given. One involved 'many', in which the sentence 'There are many stars' was presented with one picture containing six stars and the other containing two. Participants were told to click on the picture containing six stars. The other example involved an ad hoc enrichment, in which the sentence 'There is a spade' was presented with one covered picture and one picture containing a spade and a diamond. In this case, participants were instructed to click on the 'Better Picture?' option.

[2] For weak primes, the correct response was the weak picture. For strong primes, although both pictures made the sentence true, we coded the strong picture as the correct response.

There were four practice trials to familiarise participants with the task. In these trials, the sentence was either presented with a false picture and a strong picture or with a false picture and a covered picture. No feedback was given in either practice or experimental trials. The whole experiment lasted approximately 10 minutes.

3.2.4. Data treatment and analysis

The analysis was performed on the responses to target trials. Only target responses that were preceded by two correct prime responses were included in the analysis. This resulting in the removal of 35 out of 320 target responses. Of the 35, 19 were embedded targets and 16 were non-embedded targets. For the remaining target responses, we coded the enriched response as 1 and the unenriched response as 0. Note that the enriched response for embedded target trials was choosing the visible picture, whereas the enriched response for unembedded target trials was choosing the covered picture.

We fitted a logistic mixed-effect model to predict the log odds of choosing an enriched over unenriched response from fixed effects of embeddedness (embedded targets / non-embedded targets) and prime type (weak / strong). Embeddedness and prime type were deviation coded (embedded = 0.5, non-embedded = -0.5; strong = 0.5, weak = -0.5). The model contained maximal random effects structure supported by the data, which included random intercepts and slopes for subjects and random intercepts only for items. All fixed effects and their interactions were included as random slopes. Statistical analyses were carried out using R (version 3.3.3, R Core Team, 2017) with the lme4 package (Bates et al. 2015) and the lmer Test package (Kuznetsova, Brockhoff, and Christensen 2014).

3.3. Results and discussion

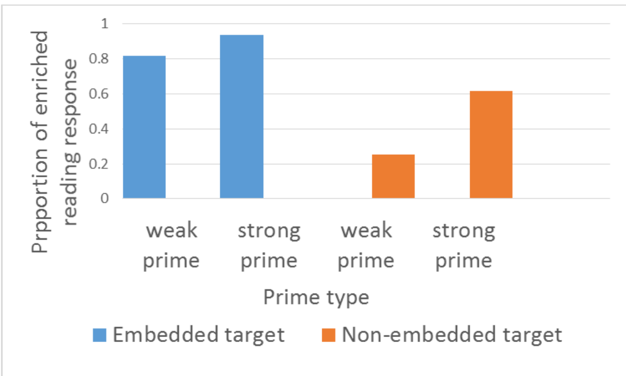

Figure 5 *The proportions of enriched responses across conditions in Experiment 1*

Figure 5 shows the proportions of enriched responses across conditions. We found a main effect of priming (β = 1.84, SE = 0.62, p = .003). However, planned comparisons on each level of prime type showed that the rate of enriched responses was significantly higher after strong primes than after weak primes only in unembedded target conditions (β = 3.48, SE = 1.36, p = .01) but not in embedded target conditions (β = 4.55, SE = 3.87, p = .24). Thus, the observed priming effect was mainly driven by the priming in the unembedded target condition. There was a main effect of embeddedness (β = 4.81, SE = 1.22, $p < .001$), suggesting that the overall rate of enriched responses was higher for embedded target trials than for unembedded target trials. The interaction between embeddedness and prime type was not significant (β = -2, SE = 1.42, p = .16).

The main effect of embeddedness in the present study is inconsistent with findings from previous research that demonstrate unembedded scalar enrichments are more robust than embedded cases (e.g. Benz & Gotzner, 2014; Geurts & Pouscoulous, 2009). However, it is difficult to read too much into this result, since the enriched response in the embedded target condition is the open card, while the enriched response in the unembedded target condition is the covered card.

Regarding whether unembedded enrichments could prime embedded enrichments, the results of this experiment are difficult to interpret. On the one hand, there is a main effect of prime type and we found no significant interaction. On the other hand, we failed to find a significant difference between Strong and Weak conditions in the embedded target condition. The main effect was driven by the significant difference between Strong and Weak trials in the unembedded target condition. This latter result is supportive of the idea that there is a shared mechanism between unembedded and embedded scalar enrichments. However, an alternative explanation for this priming effect could be given without appealing to local enrichment. Consider the items in Figure 4 again. As long as participants access the reading *On each row some of the symbols are ticks and it is not the case on each row all of the symbols are ticks*, they would choose the strong picture. This means that local enrichment is not required in deriving this reading. Enriched responses in embedded primes could be the result of global inference mechanism. Then what seems to be a local → local priming would turn out to be a global → global priming. Thus, the priming effect in unembedded target condition cannot be taken as conclusive evidence for a shared mechanism in deriving unembedded and embedded enrichment.

4. Experiment 2

In order to properly explore whether embedded and unembedded enrichments could prime each other, we conducted Experiment 2, which addressed the problems of interpreting the results of Experiment 1.

4.1. Method

4.1.1. Participants

30 participants were recruited via Prolific Academic (http://prolific.ac). All participants were native English speakers.

4.1.2. Materials, procedure

The materials were similar to Experiment 1 with one key difference, namely that for the embedded prime trials, the prime sentence was of the form *On exactly one row, some of the* *symbols are [symbol]*. As illustrated in Figure 6, in strong primes, the sentence was presented with a picture that made the literal reading true and a picture that made only the local reading true. If the participants access the local enriched reading, *On exactly one row, some but not all of the symbols are ticks*, then the only picture that made the sentence true is the 'local' picture. Since embedded enrichments in the non-monotonic environment can only be explained by local enrichment, in Experiment 2, participants who choose 'local' picture in embedded prime trials must access local enrichment.

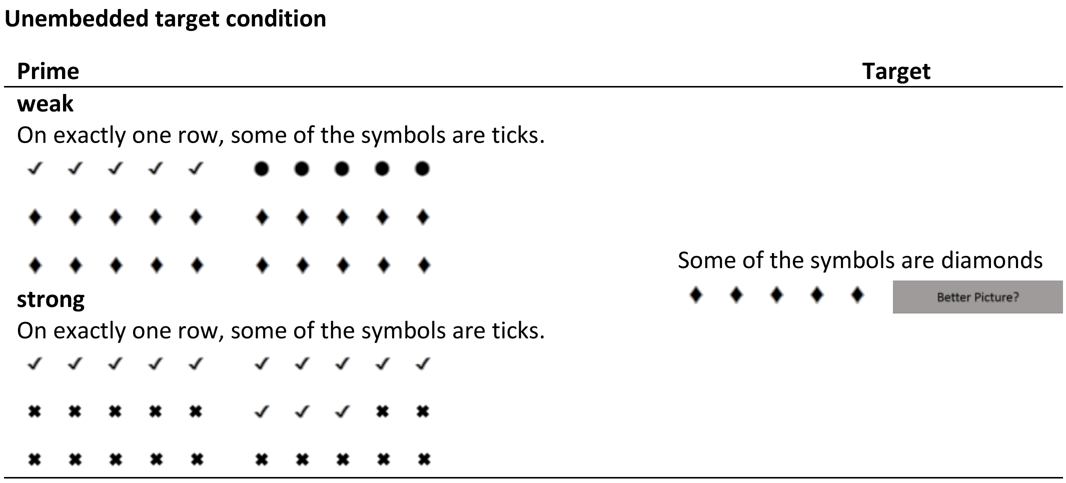

Figure 6 *Critical items for unembedded target condition in experiment 2*

As with Experiment 1, 48 filler trials were constructed. The filler sentence was of the form *All of the symbols are [symbol]* or *On exactly one row, all of the symbols are [symbol]*. Like in Experiment 1, each type of filler sentence occurred in three situations: (i) the sentence was presented with a strong picture and a 'Better Picture?', (ii) the sentence was presented with a false picture and a 'Better Picture?', and (iii) the sentence was presented with a false picture and a strong picture. All the other materials and the procedure were the same as Experiment 1.

4.1.3. Data treatment and analysis

As in Experiment 1, the analysis was performed on target responses that were preceded by two correct prime responses. This resulting in the removal of 84 out of 480 target responses. Of the 84, 24 were embedded targets and 60 were non-embedded targets. For the remaining target responses, we coded the enriched response as 1 and the unenriched response as 0.

Again we fitted a logistic mixed-effect model to predict the log odds of choosing an enriched over unenriched response from fixed effects of embeddedness (embedded / non-embedded) and prime type (weak / strong). The model contained random intercepts and slopes for subjects and random intercepts only for items. All fixed effects were included as random slopes.

4.2. Results and discussion

Figure 7 shows the proportions of enriched responses across conditions. There was a main effect of priming (β = 1.33, SE = 0.39, $p < .001$). Again, planned comparisons on each level of prime type showed that the rate of enriched responses was significantly higher after strong primes than after weak primes only in unembedded target conditions (β = 1.56, SE = 0.54, p = .004) and not in embedded target conditions (β = -1, SE = 1.71, p = .56). There was no main effect of embeddedness (β = 2.07, SE = 1.23, p = .07), and the interaction between embeddedness and prime type was not significant (β = -0.75, SE = 0.77, p = .33).

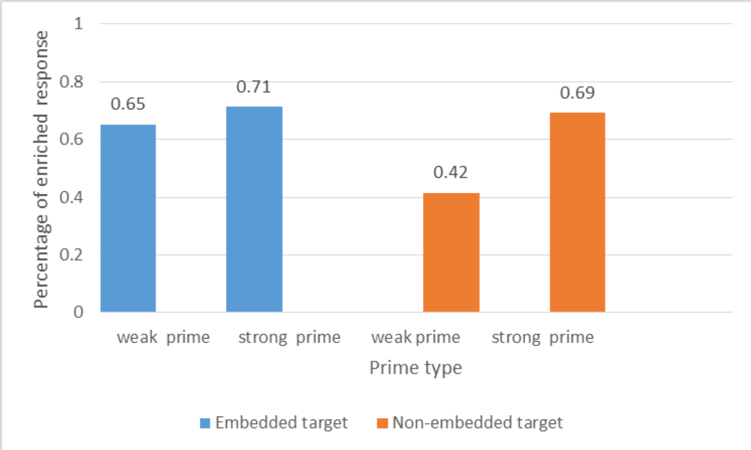

Figure 7 *The proportions of enriched responses across conditions in Experiment 2*

In this experiment, enriched responses in both embedded prime and embedded target trials could not be the product of a global enrichment. Thus, the main effect of prime types provides clear evidence that embedded and unembedded scalar implicature share a mechanism. In particular, the priming of the enriched response in the unembedded target by the embedded prime provides somewhat more direct evidence that unembedded scalar enrichments can be derived by the mechanism for local enrichment.

Overall, the main effect of prime provides support to both GT and RSA-LU accounts. In terms of discriminating between the two approaches, once again, the results are difficult to interpret, although suggestive. On the one hand, we found a priming effect in the unembedded target condition but not the embedded target condition; on the other hand, the interaction did not reach significance. It is also worth noting that the items in the embedded target condition were identical across both experiments and in both cases no effect was found in either case. As mentioned above, the RSA-LU approach predicts that, if there were an asymmetry in the priming effect, it would occur in the direction found. This is because, while embedded prime trials involve mandatory enrichment, unembedded prime trials do not. Thus the RSA approach suggests a stronger priming effect in the unembedded target condition.

5. Inverse Preference and Frequency of Local Enrichment

In this section, we will relate the results of Experiment 2 to the so-called 'Inverse Preference Effect'. Inverse preference is the phenomenon whereby a less frequent parse of a word or structure gives rise to a larger priming effect than more frequent parses (Hartsuiker, Kolk, and Huiskamp, 1999; Hartsuiker and Westenberg, 2000; Hartsuiker and Kolk, 1998;

Scheepers, 2003). For example, studies that manipulate active and passive syntactic structures find that passives, which are the less frequent construction, give rise to larger priming effects than actives (Bock, 1986). Currently favoured explanations of this effect revolve around the idea that priming itself is a result of implicit learning (Pickering and Ferreira, 2008) and that inverse preference results from error correction (Jaeger and Snider, 2013).

Inverse preference is relevant to the results in Bott & Chemla (2016), because it potentially helps to explain a surprising result in their main experiment. This is the fact that Bott & Chemla found a main effect of Within / Between, such that there were more enriched responses in the Between condition than Within, even though there was a significantly bigger effect of prime in the Within condition. This can be explained in terms of inverse preference if it is assumed that the unenriched response in prime trials is the less frequent or somehow unexpected one. This means that for Weak prime trials, there is a large priming effect for the unenriched response, causing participants to select the open picture in target trials. Bott & Chemla observe that the large priming effect in Within trials is indeed mostly due to a below baseline response in Weak trials. That is, compared to a condition where the prime was unrelated to the target in terms of scalar implicature, participants made fewer enriched responses in the Weak prime condition.

Let us now turn back to the results of Experiment 2 to consider where there might be an inverse preference effect. When we consider the unembedded target condition, it could be that because unenriched ‘some’ in Weak prime trials is unexpected, this primes the unenriched interpretation in the target. However, if the priming effect in unembedded target trials is because of below-baseline rates in weak trials, this would not explain why a similar effect is not obtained in the embedded target condition. Of course, it could be that, again, we simply failed to find the same below-baseline effect in this condition. Alternatively, if there are two mechanisms involved in scalar implicature, the literal interpretation of ‘some’ may be intermediate in its expectedness between a more frequent globally enriched reading and a less frequent locally enriched reading. This would explain the large priming effect in unembedded target trials, because the Strong primes in this condition require local enrichment and, by hypothesis, local enrichment is a less frequent response than no enrichment.

When it comes to the Embedded target condition, if global enrichment is more often used to respond to strong unembedded prime trials than local enrichment, and literal unenriched meanings are used in weak trials, then we should not expect to see such a great priming effect, because the target trials require local enrichment. This would mean that, although both global and local processes may be responsible for unembedded scalar enrichments, the global process may be the more common route.

At present, we have too little data to discriminate among these possibilities. Further studies would be required to shed light on the relation between global and local scalar enrichments in terms of their frequency. At a minimum, we would need to include an unrelated control condition here to get a better baseline.

6. Conclusion

The primary aim for this paper was to use the enrichment priming paradigm to determine whether embedded scalar enrichments and unembedded scalar enrichments involve a shared mechanism. In two experiments, we found evidence supporting a shared mechanism. In particular, Experiment 2 showed clearly that embedded prime trials, where local enrichments are mandated, lead to more unembedded scalar enrichments in targets than when only the literal meaning of ‘some’ is used in primes. This latter result in particular highlights that activation of locally enriched meanings of ‘some’ can impact on rates of unembedded scalar enrichments.

Although there are relevant differences between the RSA-LU and GT, the data in this paper does not conclusively favour one or the other. However, a twice-replicated lack of effect in the embedded target condition fits better with the Gricean picture than the Grammatical one. Again, more studies would be needed to pursue this matter further. For instance, a similar kind of study that mixes lexical triggers in a unembedded target condition might provide such a test. We leave this question open for future research.

Finally, a speculative discussion about whether the results reported in Experiment 2 might be the result of an inverse preference effect led to the suggestion that perhaps the locally enriched interpretation of ‘some’ is less frequent or more surprising than either the globally enriched or literal interpretation.

References

Bates, D., M. Maechler, B. Bolker, and S. Walker (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67 (1), pp. 1–48.

Benz, A. and N. Gotzner (2014). Embedded Implicatures Revisited: Issues with the Truth-Value Judgment Paradigm. In *Proceedings of the Formal & Experimental Pragmatics Workshop, European Summer School for Language, Logic and Information (ESSLLI)*, pp. 1–6. https://doi.org/10.3765/sp.2.4.

Bergen, L., R. Levy, and N. D. Goodman (2016). Pragmatic Reasoning through Semantic Inference. *Semantics and Pragmatics 9*, https://doi.org/10.3765/sp.9.20.

Bock, J. K. (1986). Syntactic Persistence in Language Production. *Cognitive Psychology 18 (3)*, pp. 355–387. https://doi.org/10.1016/0010-0285(86)90004-6.

Bott, L., and E. Chemla (2016). Shared and Distinct Mechanisms in Deriving Linguistic Enrichment. *Journal of Memory and Language 91*, pp. 117–140. https://doi.org/10.1016/j.jml.2016.04.004.

Carston, R. (1998). Informativeness, Relevance and Scalar Implicature. *Pragmatics And Beyond New Series*, pp. 179-238. http://discovery.ucl.ac.uk/22922/.

Carston, R. (2002) *Thoughts and Utterances : The Pragmatics of Explicit Communication*. John Wiley & Sons.

Chemla, E., and B. Spector (2011). Experimental Evidence for Embedded Scalar Implicatures. *Journal of Semantics 28 (3)*, pp. 359–400. https://doi.org/10.1093/jos/ffq023.

Chierchia, G. (2004). Scalar Implicatures, Polarity Phenomena and the Syntax/pragmatics Interface. *Structures and Beyond*, pp. 39–103.

Chierchia, G., D. Fox, and B. Spector (2012). Scalar Implicature as a Grammatical Phenomenon. In P. Portner, C. Maienborn, and K. von Heusinger (Eds.), *Semantics: An International Handbook of Natural Language Meaning*, Vol. 3 pp. 2297–2331. Berlin:

Mouton de Gruyter.

Cohen, L.J. (1971). Some remarks on Grice's views about the logical particles of natural language. In *Pragmatics of natural languages*, pp. 50-68. Dordrecht: Springer

Frank, M. C. and N. D. Goodman (2012). Predicting Pragmatic Reasoning in Language Games. *Science 336 (6084)*, pp. 998–998. https://doi.org/10.1126/science.1218633.

Gazdar, G. 1979. *Pragmatics: Presupposition, Implicature, and Logical Form*. Academic Press.

Geurts, B. (2010). *Quantity Implicatures*. Cambridge: Cambridge University Press. https://doi.org/10.1016/j.neuroimage.2009.02.016.

Geurts, B., and N. Pouscoulous (2009). Embedded Implicatures?!?. *Semantics and Pragmatics 2 (4)*, pp. 1–34. https://doi.org/10.3765/sp.2.4.

Goodman, N. D. and M. C. Frank (2016). Pragmatic Language Interpretation as Probabilistic Inference. *Trends in Cognitive Sciences 20 (11)*, pp. 818–829. https://doi.org/10.1016/j.tics.2016.08.005.

Hartsuiker, R. J. and H. H. J. Kolk (1998). Syntactic Persistence in Dutch. *Language and Speech 41 (2)*, pp. 143–84. https://doi.org/10.1177/002383099804100202.

Hartsuiker, R. J., H. H. J. Kolk, and P. Huiskamp (1999). Priming Word Order in Sentence Production. *The Quarterly Journal of Experimental Psychology 52 (1)*, pp. 129–147. https://doi.org/10.1080/713755798.

Hartsuiker, R. J., and C. Westenberg (2000). Persistence of Word Order in Written and Spoken Sentence Production. *Cognition 75*, pp. 27–39. https://doi.org/10.1016/S0010-0277(99)00080-3.

Huang, Y. T., E. Spelke, and J. Snedeker (2013). "What Exactly Do Numbers Mean?" *Language Learning and Development 9 (2)*, pp. 105–129. https://doi.org/10.1080/15475441.2012.658731.

Jaeger, T. F. and N. E. Snider (2013). Alignment as a Consequence of Expectation Adaptation: Syntactic Priming Is Affected by the Prime's Prediction Error given Both Prior and Recent Experience. *Cognition 127 (1)*. pp. 57–83. https://doi.org/10.1016/J.COGNITION.2012.10.013.

Kuznetsova, A., P. B. Brockhoff, and R. H. B. Christensen (2014). lmerTest: Tests for Random and Fixed Effects for Linear Mixed Effect Models. *R Package Version*. https://doi.org/http://CRAN.R-project.org/package=lmerTest.

Pickering, M. J. and V. S. Ferreira (2008). Structural Priming: A Critical Review. *Psychological Bulletin 134 (3)*, pp. 427–59. https://doi.org/10.1037/0033-2909.134.3.427.

Potts, C., D. L., R. Levy, and M. C. Frank (2015). Embedded Implicatures as Pragmatic Inferences under Compositional Lexical Uncertainty. *Journal of Semantics*, pp. 1–48. https://doi.org/10.1093/jos/ffv012.

Raffray, C. N., and M. J. Pickering (2010). How Do People Construct Logical Form During Language Comprehension? *Psychological Science 21 (8)*, pp. 1090–1097. https://doi.org/10.1177/0956797610375446.

Sauerland, U. (2004). Scalar Implicatures in Complex Sentences. *Linguistics and Philosophy 27*, pp. 367–391. https://doi.org/10.1023/B:LING.0000023378.71748.db.

Scheepers, C. (2003). Syntactic Priming of Relative Clause Attachments: Persistence of Structural Configuration in Sentence Production. *Cognition 89 (3)*. pp. 179–205. https://doi.org/10.1016/S0010-0277(03)00119-7.

Wilson, D. (1975). *Presuppositions and Non-Truth-Conditional Semantics*. Academic Press.