

The Path of Presupposition Projection in Processing - The Case of Conditionals

Florian Schwarz — University of Pennsylvania

Sonja Tiemann — University of Tübingen

Abstract. While presupposition projection has been explored in depth in theoretical terms, not much is known about possible cognitive processes related to it. We investigate the question of whether projection, which can be seen as a mismatch between the location of where material is introduced syntactically and the level at which it is interpreted, comes with a processing cost that requires additional processing time. In an eye tracking reading study using conditional sentences with German *wieder* ('again') in the consequent and a preceding context sentence, we varied the relative location in which the presupposition of *wieder* was supported. We argue that the pattern of results we get reflects a processing cost for projection that directly correlates with the length of the projection path as it would be measured in Discourse Representation Theory (DRT). This perspective on the data crucially relies on the representational mechanisms involved in projection, which implies that a non-representational theory such as dynamic semantics cannot offer the same account. A simple processing hypothesis based on distance from presuppositional support measured in clauses also does not suffice to explain the data. The paper thus offers experimental evidence bearing on fine-grained theoretical choices in presupposition theory. It closes by relating our results to some other recent experimental work on presupposition interpretation.

Keywords: presuppositions. presupposition projection. experimental pragmatics. DRT. dynamic semantics.

1. Background

In recent years, phenomena at the semantics-pragmatics interface have become a central area of interest in psycholinguistic experimentation. In particular, experimental work on scalar implicatures has proven to be fruitful in providing new insights in this domain. The experiment reported here is part of an effort to extend this general approach to presuppositions. The past few years have seen a growing body of work that tries to assess the interpretive properties of presuppositional content, much of it using off-line behavioral measures. Some of the issues addressed in the literature include the strength of contextual constraints imposed by a variety of presupposition triggers (Jayez and van Tiel, 2011; Amaral et al., 2011; Smith and Hall, 2011), the effects of presuppositional content on resolving ambiguities (Schwarz, 2007), and the exact nature of presuppositions in conditionals and under quantification (Chemla and Schlenker, 2009; Chemla and Bott, 2012; Romoli et al., 2011).

There also is a growing body of work trying to understand the online processing of presuppositions at a more detailed level by looking at various measures reflecting the time course of interpreting presuppositions in online interpretation. Based on the general notion that presuppositions require

some form of contextual support, previous studies have explored different experimental methods for investigating them. Schwarz (2007) found reading time effects for the part of the sentence containing *also* when the preceding sentential context did not support its presupposition. Building on this paradigm, Tiemann et al. (2011) investigated a broader range of presupposition triggers and found that unsupported presuppositions gave rise to decreased acceptability and increased reading times on the presupposition trigger itself in word-by-word self-paced reading. The increases in reading times in these studies are attributed to the clash between the context and the presupposition trigger and thus can be seen as indicative of the availability of the presuppositional content, since a mismatch can only be noticed when the presupposition has been fully computed. In a different approach, Chemla and Bott (2012) investigate reaction times for different interpretive options of presupposition triggers like *realize* under negation and report that global interpretations are faster than local ones.

The experiment reported here, building on an earlier study summarized below (Schwarz and Tiemann, 2012), follows the general approach of the reading time studies just mentioned in that it involves target sentences containing a presupposition trigger, German *wieder* ('again'), presented in context, parts of which either do or do not support the presupposition. However, we use eye tracking during reading, rather than self-paced reading, in order to allow for a more natural reading experience for the subjects and to have a more fine grained temporal resolution, which allows us to capture effects closer to their real time course. And rather than considering presupposition triggers in simple sentences without any embedding, as in most earlier studies, we consider two types of embedding environments, namely negation and conditionals. This allows us to investigate time course effects related to presupposition projection and provides a broader and more detailed picture of presupposition projection in processing and its implications for presupposition theory. In particular, we will compare a theory of presupposition projection along the lines of Discourse Representation Theory (DRT) and one provided in the framework of Dynamic Semantics. Both are designed to capture the integration of sentence meaning and discourse context, including presuppositions. Yet, plausible predictions about the processing of embedded presuppositions based on these two theories are rather different, as we discuss in detail below. In closing, we will briefly discuss the highly relevant findings of (Chemla and Bott, 2012) in relation to our results in more detail in the general discussion section.

1.1. Theoretical Perspectives on Projection in Processing

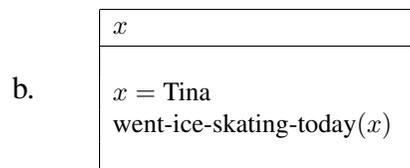
One of the core properties of presuppositions is that they project out of a variety of embedded contexts. To illustrate briefly: whereas only the declarative in (1-a) conveys that Tina went ice-skating today - which we take to be the asserted or proffered content of the sentence, which is directly affected by the embedding operators -, the presupposition that Tina went ice-skating before remains constant across all four variations of the sentence:

- (1) a. Tina went ice-skating again today.
 b. It is just not true that Tina went ice-skating again today.
 c. If the weather was nice, then Tina went ice-skating again today.
 d. Did Tina go ice-skating again today?
- (2) **Presupposition:** Tina went ice-skating before.

Descriptively speaking, projection involves a discrepancy between the syntactic location in which a presupposition is triggered and the level where it is interpreted. Intuitively, disentangling this mismatch could plausibly lead to difficulties in sentence processing. However, since projection is such a core property of presuppositions, it also seems possible that the underlying mechanism proceeds in an automated way that does not incur any extra effort. The question of which of these two characterizations turns out to be correct is not only interesting from a processing point of view, it also bears on ongoing theoretical controversies, as one of the major challenges for compositional semantic theories is to formulate theoretical mechanisms that account for the projection behavior of presuppositions. In the following, we will focus on two ‘classic’ semantic theories which integrate sentence meaning and discourse context in a dynamic manner. Although Dynamic Semantics (Heim (1983), Groenendijk and Stokhof (1990) , Chierchia (1995), among many others) and DRT (Kamp (1981), van der Sandt (1992), van der Sandt and Geurts (1991), Geurts (1999)) make similar predictions about the projection behavior of presuppositions, their technical implementations differ substantially in ways that suggest different processing predictions. The differences are mainly due to the fact that a DRT analysis of presuppositions involves operations on representations of discourse structures (DRS), whereas Dynamic Semantics accounts for them within a framework where sentences directly update the information encoded by the context in a non-representational way. In particular, while it seems plausible in DRT to assume that presupposition projection is an effortful process with additional cognitive steps involved, Dynamic Semantics provides no grounds for expecting such additional efforts . The implementations of these two theories and their potential impact on processing are laid out in more detail below.

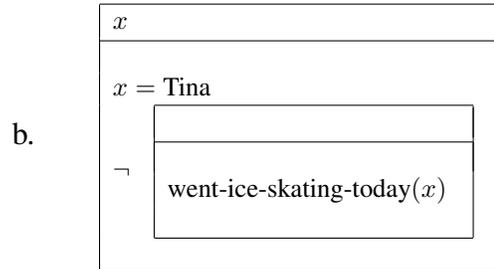
Discourse Representation Theory (DRT) is a representational theory of meaning in discourse, in which a Discourse Representation Structure (DRS) is continuously constructed as the discourse unfolds. The DRS makes it possible to keep track of already established information in the discourse, and to relate newly added information to it. Discourse representations can be conveniently represented using boxes, which consist of a header and a main body. The header contains a list of discourse referents, while the main body contains descriptive conditions, i.e., it lists predicates which hold of certain of the discourse referents. A simple illustration for (3-a) is given in (3-b).

- (3) a. Tina went ice-skating today.



Embedding operators add additional embedded DRSs, as illustrated in (4) for negation.

(4) a. Tina didn't go ice-skating today.



DRT sees presuppositions as anaphoric elements, which have to be linked to a discourse antecedent. More precisely, presuppositions that are introduced in embedded contexts follow a pre-defined projection path through the complex DRS-structure, in search for the closest possible antecedent. According to this perspective, interpreting a presupposition involves a mechanism that performs operations on DRSs to connect it to its antecedent. Consider the following pseudo-English versions of the types of German constructions investigated in Schwarz and Tiemann (2012) and in the experiment below:

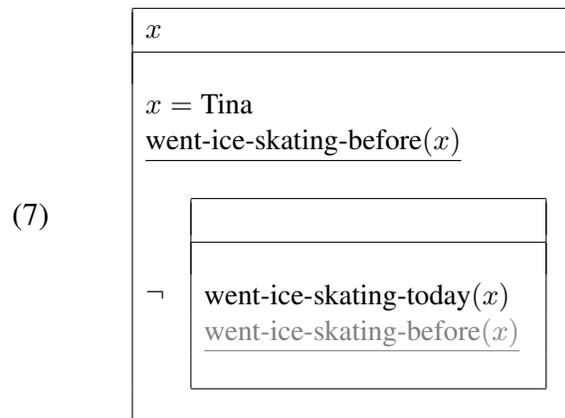
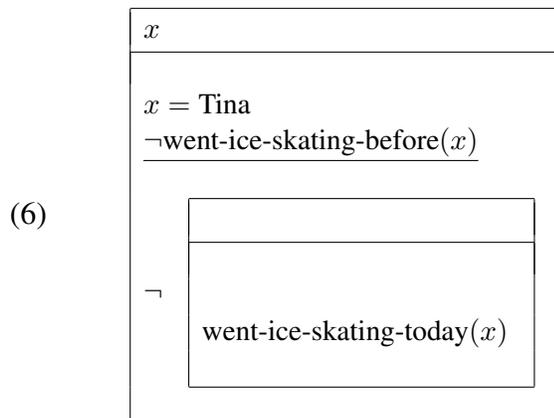
(5) a. Tina AGAIN NOT went ice-skating today.

(≈ Once again, Tina did not go ice-skating today.)

b. Tina NOT AGAIN went ice-skating today.

(≈ It is not true that Tina went ice-skating again today.)

In (5-a) the presupposition trigger *again* appears outside of the scope of negation, triggering the presupposition that there was a salient earlier time at which Tina did not go ice-skating, which is introduced at the global level. A DRT representation of (5-a) is given in (6).



In (5-b) on the other hand, the presupposition is triggered in the scope of negation, yielding the

representation in (7). Initially, the presupposition that Tina went ice-skating before is introduced inside of the sub-DRS introduced by negation. In order to be resolved to a DRS-condition introduced in the preceding discourse, the (underlined) presupposition has to move to the global DRS. (We use light-gray font to indicate locations at which a presupposition was represented at earlier stages of the DRS construction process.) As can easily be seen when comparing these two cases, interpreting a presupposition that was triggered in an embedded construction involves an extra step compared to a sentence like (5-a), where the presupposition is generated in the global DRS to begin with. If we assume that the added step in (5-b) corresponds with extra processing time, this predicts that interpreting embedded presuppositions will be slower than interpreting unembedded ones.

Dynamic Semantics characterizes the meaning of a sentence in terms of its context change potential, i.e., by characterizing what impact it would have in any given context. A context, in the simplest terms, can be seen as a set of possible worlds which is updated by every sentence uttered. This set corresponds to the information that counts as established in the discourse. Presuppositions constrain the update process. In particular, context update based on a sentence can only proceed if all its presuppositions are true in all possible worlds that make up the context. A simple context update without presuppositions is illustrated in (8) where c is the original context, and c' the resulting context after updating c with a sentence S .

$$(8) \quad c + S = c'$$

Adding presuppositions into the mix, a given context only *admits* a sentence if it entails, or *satisfies* all the sentence's presuppositions. This can be expressed as a definedness condition:

- (9) a. Tina went ice-skating again today.
 b. $c + \text{Tina went ice skating again today} = c'$
- defined iff $c + \text{Tina went ice-skating before} = c$
 - if defined, c' will only contain those worlds in c in which Tina went ice-skating today

The context change potential of a complex sentence is determined by the context change potential of its parts. The context change potential of a negated sentence is illustrated in (10).

$$(10) \quad c + \text{Not } S = c - (c + S)$$

What is crucial to note for our purposes is that the formulation in (10) requires the presuppositions of S to be entailed by c , just as was the case for non-negated sentences. An illustration for the example in (5-b) from above is given in (11) where (11-a) is the sentence that negation applies to, (11-b) its presupposition (PSP), and (11-c) its context change potential.

(5-b) Tina NOT AGAIN went ice-skating today.

- (11) a. S: Tina went ice-skating again today.
b. PSP: There is a (salient) earlier time where Tina went ice-skating.
c. $c - (c+ \text{Tina went ice-skating today.})$
(defined iff $c+ \text{Tina went ice-skating before} = c$)

As shown in (11) the initial context c has to admit S in order to compute $c + \neg S$. The same holds true for the version of the sentence where the presupposition is triggered outside of the scope of negation (although with this particular example, the presupposition is different due to the presence of negation):

(5-a) Tina AGAIN NOT went ice-skating today.

- (12) a. S: Tina went ice-skating today.
b. PSP: There is a (salient) earlier time where Tina did not go ice-skating.
c. $c - (c+ \text{Tina went ice-skating today.})$
(defined iff $c+ \text{Tina did not go ice-skating before} = c$)

Evaluating presuppositions in embedded and unembedded environments thus proceeds in exactly the same way in this framework. In both cases, the presupposition of the sentence is checked against the global context. Therefore, there is no reason to expect the interpretation of presupposed content in cases where the presupposition projects to be cognitively more demanding than in cases where it is introduced globally from the start. This contrasts, of course, with the prediction of the DRT account spelled out above, and the experiment reported here aims to test precisely this divergence in processing predictions suggested by the two theories of presupposition projection.

1.2. Prior Work on Processing Projection out of the Scope of Negation

A first attempt at testing the predictions laid out above is reported in Schwarz and Tiemann (2012), where we present results from a previous reading time study using eye tracking. The experiment discussed there manipulated the embedding of German *wieder* ('again'). In German, *wieder* and

negation can appear adjacently in either order, as shown in the variations of the target sentence in (14). As was already seen in the discussion of (5) above, the two different word orders yield different presuppositions: while the presupposition of (a) is that there is a salient earlier event of Tina going ice-skating, (b) presupposes that there exists a salient earlier event at which Tina did not go ice-skating. We presented the target sentences in the context of a preceding sentence which was varied so as to manipulate whether the presupposition of the target sentence was supported (C1 for (a), C2 for (b)) or not (C2 for (a), C1 for (b)).

(13) **C1:** *Tina went ice skating for the first time last week with Karl. The weather was beautiful, and they had a great time.*

C2: *Tina wanted to go ice skating for the first time last week with Karl. But the weather was miserable and they gave up on their plan.*

(14) Dieses Wochenende war Tina {(a) **nicht wieder** / (b) **wieder nicht**} Schlittschuhlaufen,
 this weekend was Tina {(a) **not again** / (b) **again not**} ice-skating
 weil das Wetter so schlecht war.
 because the weather so bad was

Looking at reading times on the verb after the *{nicht wieder / wieder nicht}* region, we observed an interaction of the two factors **Felicity** and **Firstword** for a variety of reading measures, including very early ones (first fixation, first pass time, regression path duration, and total time). Reading times were increased in the infelicitous condition when the first word was *wieder*, but when *nicht* was the first word, the felicitous and the infelicitous condition did not differ in terms of reading times.

In order to test whether subjects perceived the infelicitous *nicht wieder* condition as infelicitous after all, or whether they resorted to a local interpretation of the presupposition with respect to negation, which would render a felicitous discourse, a follow-up acceptability rating study was carried out. It revealed that subjects rated the overall discourse in the infelicitous condition as significantly less acceptable than in the felicitous condition, for both word orders. This result indicates that while an effect of felicity was absent in the *nicht wieder* condition in the reading time data, subjects perceived a difference in the off-line acceptability judgment task, ruling out the possibility that they interpreted the presupposition in the *nicht wieder* condition locally with respect to negation. The interpretation we propose in Schwarz and Tiemann (2012) is that projection out of an embedded environment takes time, and that this is most compatible with a DRT analysis, which posits explicit and complex operations on levels of representation in the computation of a global interpretation of a presupposition introduced in an embedded environment, as illustrated above. The results of this experiment thus provide initial evidence for the hypothesis that discourse representations such as DRSs in general, and the corresponding operations on them that presupposition projection involves in particular, have real cognitive correlates.

Condition	Context	Location	Firstword
a	I	local	<i>wieder</i>
b	I	global	<i>nicht</i>
c	II	global	<i>wieder</i>
d	II	local	<i>nicht</i>

Table 1: Overview of Conditions and Factors

2. Experiment: Projection in Conditionals

To further test the hypothesis that presupposition projection is an effortful process, due to the complexity introduced by the required operations on discourse representations, the materials from Schwarz and Tiemann (2012) were modified so as to include a conditional, which provides an additional layer of embedding. Another change from the previous experiment was that all of the discourses had a felicitous interpretation, though with variation in the location where support for the presupposition was introduced. This provides a broader perspective on how projection is realized in processing, and allows us to compare different processing accounts of projection related effects in even more detail.

2.1. Methods

Design & Materials Our items were created according to the following pattern: the presupposition of the target sentence was either supported by the global context sentence or by the local *if*-clause. As before, we manipulated the order of *wieder* and negation in the target sentence, which in this case was the consequent of the conditional. This yielded a 2×2 design, which could be characterized by pairs of factors in various ways. Table 1 provides an overview of the different possible groupings based on what pairs of factors we consider.

- (15) Tina war letzte Woche $\{(\mathbf{I}) \emptyset / (\mathbf{II}) \text{ nicht}\}$ Schlittschuhlaufen. Wenn sie gestern $\{(\mathbf{I})$
Tina was last week $\emptyset / \text{not}\}$ ice-skating. If she yesterday
nicht / $(\mathbf{II}) \emptyset\}$ Schlittschuhlaufen war, dann...
not / \emptyset ice-skating was, then...
- (16) ...geht sie heute bestimmt $\{\text{nicht wieder} / \text{wieder nicht}\}$ Schlittschuhlaufen.
...goes she today certainly *not again / again not* ice-skating.

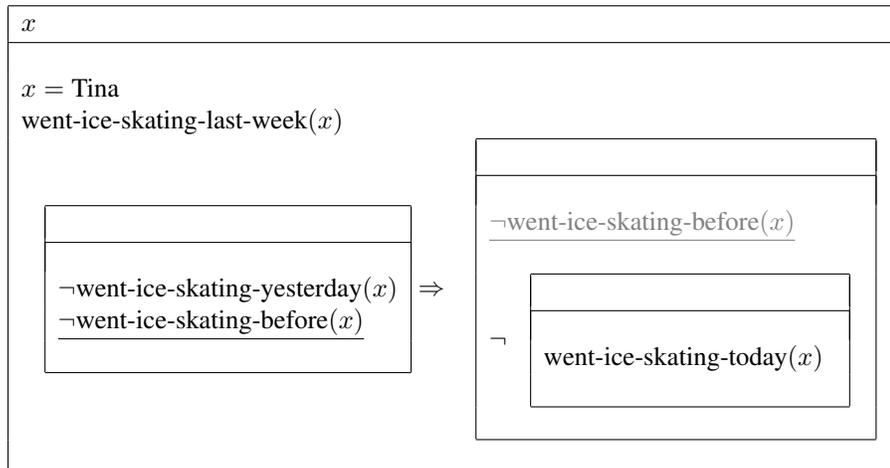
Participants and Procedure 24 items with variations of the illustrated pattern were created, each with 4 versions for the four conditions. 32 native speakers of German from the University of Tübingen community participated in the experiment. Subjects were split into 4 groups, where each subject saw 6 of the sentences per condition, providing us with a balanced number of data points from all conditions for each item and subject. There were 50 filler sentences from other, unrelated experiments. Subjects read all sentences on a computer screen while we recorded their eye movements with an EyeLink 1000 eye tracker. Half of the items were followed by a simple yes/no question to ensure that subjects were reading the materials for comprehension.

2.2. Predictions

Before turning to the results, let us spell out the predictions of the two accounts under consideration in detail. To do so, we have to introduce the way they handle conditionals in general, and presupposition projection in conditionals in particular.

DRT Predictions Operators such as conditionals introduce new embedded DRSs into the structure in DRT. Depending on where the presupposition is introduced and where a suitable antecedent can be found, different projection path lengths come about. Conditionals introduce two sub-DRSs, connected by an arrow. For presuppositions introduced in the consequent, the first place to look for support for the presupposition is the antecedent of the conditional, followed by higher levels. Let us take a look at the DRSs for our four conditions. In condition a, the presupposition that there is an earlier time where Tina did not go ice-skating is introduced in the top-level box of the consequent and satisfied locally in the antecedent of the conditional. Resolution of the presupposition then only involves one step, yielding a projection path length of 1.

(17) **Condition a**

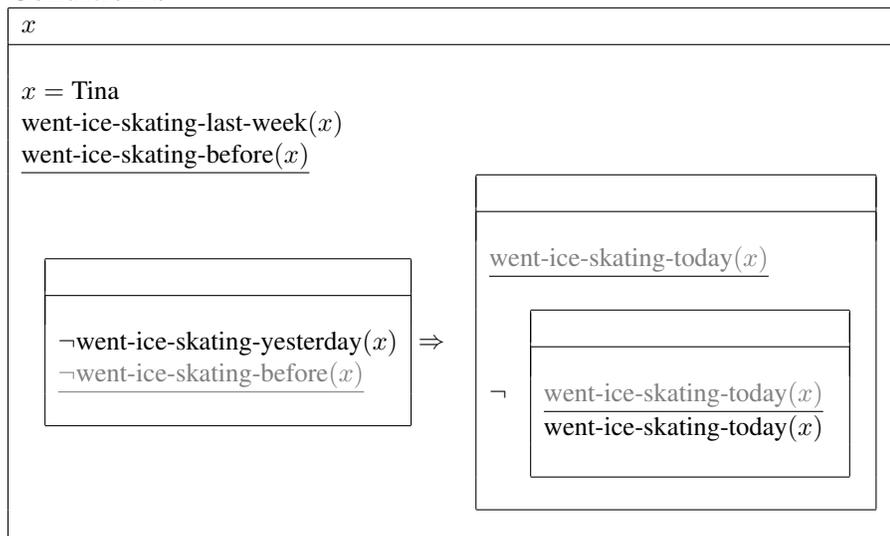


Context: I
Location: local
Firstword: wieder

Projection Path Length = 1

In condition b on the other hand, the presupposition that Tina did go ice-skating before is introduced inside of the box introduced by negation in the consequent and resolved at the top-most level, based on the information introduced through the context sentence. Given standard DRT-assumptions about the path followed in searching for an antecedent, this involves three steps.

(18) **Condition b**

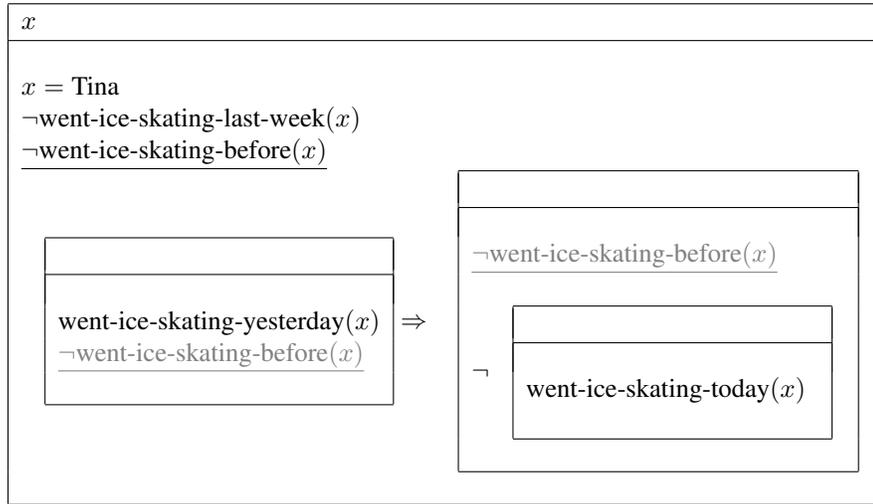


Context: I
Location: global
Firstword: nicht

Projection Path Length = 3

Turning to condition c, we again have the presupposition of the *wieder nicht* order, as in a, but this time, given Context II, it is resolved by the context sentence at the top-most level, thus rendering a projection path length of 2.

(19) **Condition c**

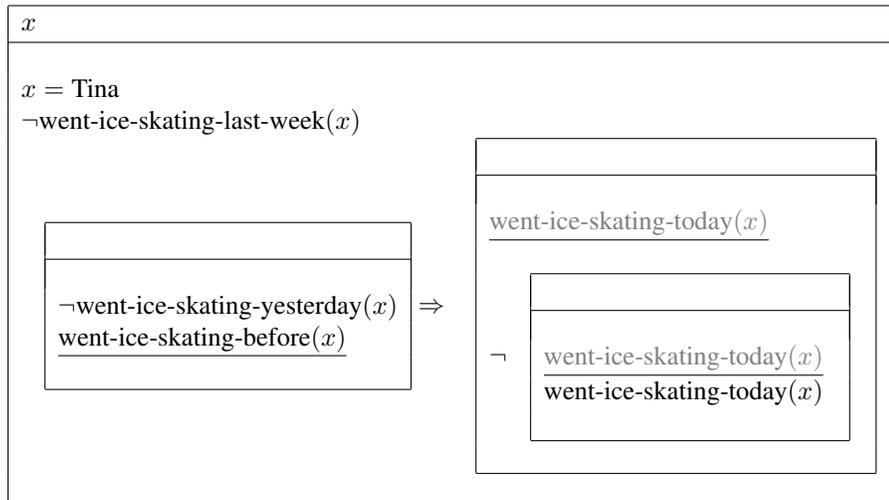


Context: II
Location: global
Firstword: wieder

Projection Path Length = 2

Finally, condition d has the same presupposition as b, introduced in the scope of negation inside of the consequent, but this time it is resolved locally in the antecedent of the conditional. This also yields a projection path length of 2.

(20) **Condition d**



Context: II
Location: local
Firstword: nicht

Projection Path Length = 2

Looking at the distribution of the projection path lengths, we see that an interesting pattern emerges. For Context I, we get a difference based on the location where support for the presupposition is introduced (corresponding to a difference between the *nicht wieder* and *wieder nicht* conditions), with a longer projection path for the global condition (3) than for the local condition (1). Con-

text II, in contrast, does not give rise to such a difference, as both the global and local conditions (corresponding to *wieder nicht* and *nicht wieder* correspondingly) have a projection path length of 2. This comes about because of the additional embedding introduced by negation. For the *nicht wieder* condition, this results in two steps, first out of the negation box into the main consequent box, then to the box for the antecedent of the conditional. For the *wieder nicht* condition, on the other hand, there also are two steps, but different ones: from starting out in the top-most box of the consequent, there's one step to the antecedent, and another to the global level. Hypothesizing that longer projection path lengths correspond to cognitive efforts that are reflected in increased reading times, this projection path distribution of the DRT account leads us to expect an interaction between **Context** and **Location**, with a difference between the a and b conditions, but no difference between c and d, as illustrated in the left panel of Figure 1. (a main effect of **Location** is also predicted, though it is dominated by the interaction).

Dynamic Semantics Predictions A simple view of conditionals in Dynamic Semantics assumes the equivalent of logical implication as their meaning contribution, i.e., *if p, then q* is assumed to be true unless *p* is true and *q* is false. To do this dynamically, and to account for presupposition projection, this is done in a slightly more round-about way, namely by removing all those worlds from the context set that remain when we add *p* to *c* and then subtract the combination of *c*, *p*, and *q* from that, thus removing precisely those worlds from *c* where *p* is true and *q* is false. The formal formulation of the context change potential of a conditional is given in (21-a).

- (21) a. $c + \text{If } p, q = c - ((c + p) - ((c + p) + q))$
 b. defined iff $(c + p) + PSP_q = (c + p)$

As far as presuppositions are concerned, and specifically the ones introduced in the consequent as in our case, what is important is that as in the case of negation, the presuppositions of *q* have to be met in the context that it's added to, but in this case, this context turns out to be the combination of *c + p*. That is, the context update for a conditional can only be performed if the global context updated with the antecedent of the conditional entails the presuppositions introduced in the consequent of the conditional. This means, however, that it should not make a difference whether the presupposition is supported by the global context (the preceding sentence) or the local context (here: antecedent of the conditional). The fact that this holds indeed is a crucial piece of the Dynamic Semantic account of presupposition projection, as this allows presuppositional support to be introduced a) in the global context, b) in the antecedent of the conditional, or c) by both combined. As far as our experimental manipulation is concerned, however, it also means that on the basis of the semantics alone, we can not differentiate between these different possible sources of support, and thus cannot come up with different processing predictions w.r.t. relative processing effort involved in these cases.

If pressed to squeeze some kind of potential prediction out of the dynamic account, one might wonder about negation as a factor. It might be reasonable to hypothesize that a negated presupposition is harder to process than an unnegated presupposition, both on intuitive grounds, and based on the more complicated update steps involved. If we take r in (22) to be the unnegated version of the presupposition in our materials, the underlying mechanisms to check the presuppositions of the *again not* and *not again* variants would involve (23-a) and (23-b) respectively.

(22) $r =$ Tina had been ice-skating before

(23) a. $PSP_q = \neg r :$ $c + \neg r = c?$ \approx $c - (c + r) = c?$
 b. $PSP_q = r :$ $c + r = c?$

Hypothesizing that negation adds extra complexity would then suggest that the *again not* sentences should be harder to process than the *not again* sentences, based on the presupposition evaluation process. However, the point about the evaluation of negation requiring extra resources presumably should be applied to the processing of non-presupposed content as well. In our case, this would suggest that negation in the antecedent is relevant as well, which would increase the the relative difficulty of the sentences in Context I, but not Context II, as illustrated for our four conditions below.

	Context I		Context II
(24)	a) $(c' + \neg q) + \neg \mathbf{PSP}_r = c' + \neg q$	c)	$(c' + q) + \neg \mathbf{PSP}_r = c' + q$
	b) $(c' + \neg q) + \mathbf{PSP}_r = c' + \neg q$	d)	$(c' + q) + \mathbf{PSP}_r = c' + q$

Factoring negation into the equation along these lines, a dynamic account thus also gives rise to the prediction of a **Location** \times **Context** interaction. However, this interaction is crucially different from the one predicted by a DRT account in that the expected simple effect of **Location** in Context I is in the opposite direction and a simple effect of **Location** is predicted in Context II (see middle panel of Figure 1).

A Processing Hypothesis However, there is at least one further perspective on the experimental materials that one could consider, based on the notion that *again* is a so-called anaphoric presupposition trigger, as it relates directly to previous material in the discourse. Assuming that some form of anaphoric resolution is involved in interpreting the presupposition introduced by *again*, one might then entertain as a plausible hypothesis the idea that cases with closer antecedents are easier to process than cases with ones that are further away. The issue arises of how exactly to quantify distance from the antecedent here, but note that most, if not all, measures that one could consider will not differentiate between the *wieder nicht* and *nicht wieder* orders, since there is only

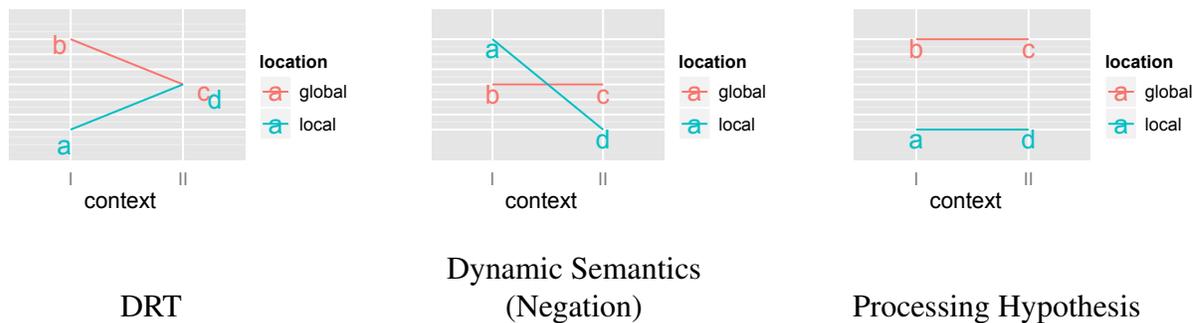


Figure 1: Processing Predictions

one extra adverb intervening between the presupposition trigger and its antecedent in the former case. If we count the distance in terms of clauses, for example, the two will come out on par. What we expect to matter, on such a view, then, is simply whether the antecedent for the presupposition is introduced in the global context sentence or in the local antecedent of the conditional. This gives rise to a straightforward processing prediction, namely a main effect of **Location**, with no differences based on **Context** (or **Firstword**). The set of predictions from the three perspectives considered are summarized in Figure 1.

2.3. Results

The primary focus in our analysis were the reading times on the verb following the *{wieder nicht}* sequence in the consequent of the conditional. Since the presupposition of the clause relies on the verb, this is the point at which the presupposition becomes fully explicit. Moreover, effects for reading measures on the verb have already been reported in Schwarz and Tiemann (2012). Standard reading measures were calculated (Rayner, 1998), namely first fixation duration, which measures the length of the very first fixation on the region of interest (here the verb); go-past time, which here is taken to measure the sum of all fixations on the region of interest prior to any fixations to the right of this region (but not including the time of regressive fixations); first pass time, which includes all fixations on the region when it is looked at the first time, up until leaving the region (to either the left or right); total duration, which sums all the fixations on the region of interest, no matter when they occur; regression path duration, which measures all fixations from first entering the region to first leaving it to the right (including all potential regressive fixations; this is sometimes also referred to as go past time); and first pass regression proportion, which is the proportion of regressive eye movements following the first time of entering the region.

All analyses used mixed-effect models with subjects and items as random effects, using the *lmer* function of the *lme4* package in *R* (Bates, 2005), with p-values calculated using MCMC estimates for significance (Baayen et al., 2008). Given recent arguments by Barr et al. (in press) that maximal

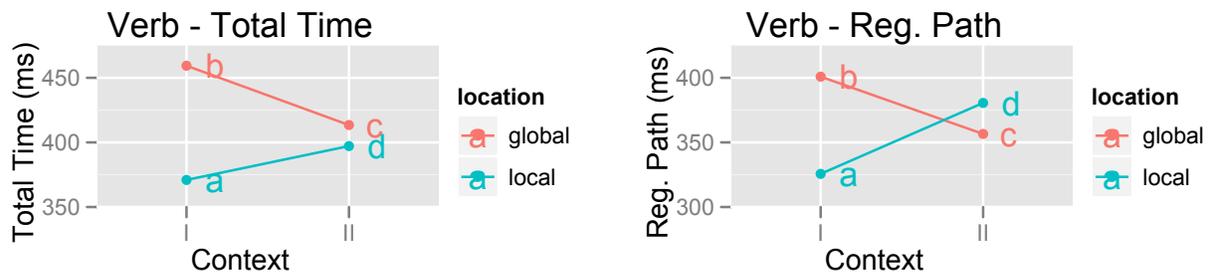


Figure 2: Summary of Reading Time on the verb in ms for Total Reading Time and Regression Path Duration

random effect structures should be used when possible, we generally computed models with the maximal random effect structure that would converge, with random effect slopes for each factor (and an interaction where applicable). We report t -values for these models as well.¹ Significant effects were found for regression path duration and total reading times. Re-reading times also gave rise to some significant effects, but the number of data points was very small for these, so we focus on the former two in the presentation of our results. Trials where any blinks occurred while looking at the region of interest were removed prior to analysis. Data points that deviated by more than two standard deviations from the mean of their condition were excluded from the analysis.

The mean reading times for the two measures to be discussed here are summarized in Figure 2. The main set of statistical analyses were carried out using **Location** and **Context** as factors, as this provides the most informative perspective on the predictions of the different accounts discussed above. The central result is that there is a significant difference between conditions a and b, but no reliable difference between c and d, i.e., the local vs. global distinction mattered in Context I (where *wieder nicht* was paired with local and *nicht wieder* with global), but not in Context II (with the reverse pairing). Statistically, this was supported by a significant interaction between context and location.

For Regression Path Duration, the model with random intercepts for subjects and items yielded a significant interaction $t = 2.95, p < 0.01$ (the model with a full random effect structure including the interaction did not converge; the maximal models that did converge, with random effect slopes for subjects for just one of the factors, yielded comparable t -values (all greater than 2.9)). There also were significant main effects of **Location** ($t = 2.11, p < 0.05$) and **Context** ($t = 3.35, p < 0.001$), but these were dominated by the cross-over interaction. The interaction was driven by the difference between conditions a and b: analyzed as a simple effect with only random intercepts, this yielded $t = 3.06, p < .01$ (including a random effect slope for **Location** yielded $t = 2.87$). The difference between conditions c and d, on the other hand, was not significant ($t \leq 1$).

¹MCMC simulations are not yet implemented for complex random effect structures in lme4; given the size of our data set, t -values greater than 2 roughly correspond to significance at the conventional $\alpha = .05$ level.

Essentially the same pattern emerged for total time measures. The interaction based on a model with random intercepts yielded $t = 2.25, p < .05$ (a model with random slopes for both factors and the interaction yielded $t = 2.18$). There again were significant main effects of **Location** ($t = 3.22, p < 0.01$) and **Context** ($t = 2.14, p < 0.05$), which were dominated by the interaction. The interaction was again driven by the difference between conditions a and b ($t = 3.93, p < .001$; for full model, $t = 3.20$), as opposed to the lack of a difference between conditions c and d ($t < 1$). In addition, there were simple effects of **Location** for both orders of *nicht* and *wieder*, i.e., significant differences between conditions a and c ($t = 2.20, p < 0.05$) and b and d ($t = 2.76, p < 0.01$).

Comparable results emerged for these two reading time measures when looking at the *{wieder nicht}* region, though we won't discuss these here in detail for reasons of space. Altogether, the pattern of the results very closely resembles the predictions of the DRT account as laid out above. The greatest and clearest difference in reading times is found for the two conditions with the greatest difference in projection path length, namely a and b, whereas no difference is found between c and d, counter to what we would expect based on the processing hypothesis. Furthermore, for total reading times, we also find differences based on the level of the **Location** factor for both *wieder nicht* and *nicht wieder*, with longer reading times in the global conditions. These results correspond surprisingly well to the pattern expected based on DRT projection path length. Based on this observation, we conducted an additional follow-up analysis, in which we tried to model the data using only projection path length as a (numerical) predictor. Both for regression path duration and total time, this yielded a significant effect of projection path length ($t's > 2.8$). Based on model comparisons with the interaction analysis above, the resulting fit was not significantly different for the two analyses. In other words, projection path length alone was as good a predictor of reading times as the **Location** × **Context** interaction.

3. General Discussion

The results for total reading time and regression path duration mirror the predictions of the DRT hypothesis based on projection path length remarkably well, and indeed the data can be modeled adequately by just considering projection path length as a predictor. This lends strong support to the idea that presupposition projection involves representational complexity of some kind, as on the DRT account, and that representational complexity has real cognitive correlates. Interestingly, no effects of distance from the antecedent in terms of anaphora processing were found, i.e., whether presuppositional support was introduced at the level of the global context sentence or in the antecedent of the conditional did not matter if one kept the overall projection path length constant by switching from the *wieder nicht* (c) condition to the *nicht wieder* condition (d). What seems to matter in terms of processing effort simply is the length of the projection path overall, not the absolute location at which support for a presupposition is introduced.

We thus have found further evidence supporting the notion, first suggested in Schwarz and Tie-mann (2012), that presupposition projection takes time. This is unexpected on accounts that model presupposition projection in purely semantic terms, such as Dynamic Semantics, as they cannot

readily distinguish between varying locations of presuppositional support. While we are not in a position to evaluate other, more recent, accounts of presupposition projection in detail here, it is worth noting that as far as projection is concerned, proposals such as Schlenker (2009) seem to be on par with Dynamic Semantics with respect to the present data in that they, too, evaluate presuppositions relative to the entire preceding discourse and are thus not able to differentiate different locations of presuppositional support. More generally, pragmatic theories that derive presuppositions via some form of pragmatic reasoning, which would seem to predict processing delays based on results from the literature on scalar implicatures, do not predict differences corresponding to the differences in DRT projection path length that are reflected in our experimental data.

In conclusion, let us note that it will be particularly interesting to compare our results in detail to those reported in Chemla and Bott (2012). These authors measured reaction times for local and global accommodation readings under negation for factive verbs such as *realize*, and find global response times to be faster than local ones. On the face of it, this seems directly contradictory to our results. However, note that in addition to looking at another type of presupposition trigger, these authors consider cases of accommodation, rather than resolution in the immediate discourse context. Furthermore, their ‘local accommodation’ interpretations are strictly local (in their case, within the scope of negation), whereas our ‘local’ resolution here is only local relatively speaking (i.e., more local than the global one). In fact, the characterization of presupposition resolution and accommodation in DRT, as first suggested by van der Sandt (1992), is very much in congruence with the data from these two studies, as it involves two phases for presupposition interpretation: in the first phase, the presupposition is passed up the projection path until it finds an antecedent; if this doesn’t succeed, the second phase consists of a search back down the projection path to find a suitable location for accommodation. Based on this, we would precisely expect local resolution to be faster than global resolution, but local accommodation to be slower than global resolution or accommodation. Needless to say, these issues have to be investigated in more detail, and a further direct comparison between these various options for interpreting presuppositions will be crucial. But while much more work is in order, we hope to have made progress on understanding presuppositions in online processing with the present contribution.

References

- Amaral, P., C. Cummins, and N. Katsos (2011). Experimental evidence on the distinction between foregrounded and backgrounded meaning. In *Proceedings from the Workshop on Projective Meaning at ELLI 2011*.
- Baayen, R. H., D. J. Davidson, and D. M. Bates (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59, 390–412.
- Bates, D. M. (2005). Fitting linear mixed models in r. *R News* 5, 27–30.
- Chemla, E. and L. Bott (2012). Processing presuppositions: dynamic semantics vs pragmatic enrichment. *Language and Cognitive Processes* DOI: 10.1080/01690965.2011.615221.

- Chemla, E. and P. Schlenker (2009). Incremental vs. symmetric accounts of presupposition projection: An experimental approach.
- Chierchia, G. (1995). *Dynamics of meaning. Anaphora, Presupposition and the Theory of Grammar*. Chicago and London: University of Chicago Press.
- Geurts, B. (1999). *Presuppositions and Pronouns*. Amsterdam; New York: Elsevier.
- Groenendijk, J. and M. Stokhof (1990). Dynamic montague grammar. In L. Kálmán and L. Pólos (Eds.), *Papers from the Second Symposium on Logic and Language*, Budapest, pp. 3–48. Akadémiai Kiadó.
- Heim, I. (1983). On the projection problem for presuppositions. In *Proceedings of the Second West Coast Conference on Formal Linguistics*, pp. 114–125.
- Jayez, J. and B. van Tiel (2011). Is ‘only’ special? Presentation at the XPRAG Workshop, Pisa, Italy.
- Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Janssen, and M. Stokhof (Eds.), *Formal Methods in the Study of Language: Proceedings of the Third Amsterdam Colloquium*, Volume I, Amsterdam, pp. 227–321. Mathematical Center.
- Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological 124*, 372–422.
- Romoli, J., Y. Sudo, and J. Snedeker (2011). An experimental investigation of presupposition projection in conditional sentences. *Proceedings of SALT 21*, 592–608.
- Schlenker, P. (2009). Local contexts. *Semantics & Pragmatics 3*, 1–78.
- Schwarz, F. (2007). Processing presupposed content. *Journal of Semantics 24*, 373–416.
- Schwarz, F. and S. Tiemann (2012). Presupposition processing - the case of german wieder. In M. Aloni, V. Kimmelman, F. Roelofsen, G. W. Sassoon, K. Schulz, and M. Westera (Eds.), *Logic, Language and Meaning: 18th Amsterdam Colloquium*, pp. 200–209.
- Smith, E. A. and K. Hall (2011). Projection diversity: Experimental evidence. In *Proceedings from the Workshop on Projective Meaning at ESLLI 2011*.
- Tiemann, S., M. Schmid, N. Bade, B. Rolke, I. Hertrich, H. Ackermann, J. Knapp, and S. Beck (2011). Psycholinguistic evidence for presuppositions: On-line and off-line data. In I. Reich, E. Horch, and D. Pauly (Eds.), *Sinn & Bedeutung 15 Proceedings of the 2010 Annual Conference of the Gesellschaft für Semantik*, pp. 581–597.
- van der Sandt, R. (1992). Presupposition projection as anaphora. *Journal of Semantics 9*, 333–377.
- van der Sandt, R. and B. Geurts (1991). Presupposition, anaphora, and lexical content. In *Text Understanding in LILOG*, pp. 259–296.