How and why conventional implicatures project

Noortje J. Venhuizen
University of Groningen

Johan Bos
University of Groningen

Petra Hendriks
University of Groningen

Harm Brouwer
Saarland University

Abstract Conventional Implicatures (CIs; in the sense of Potts 2005) are part of a larger class of projection phenomena. These phenomena also include presuppositions and anaphora, and can be described as content that is not *at-issue* (see Simons, Tonhauser, Beaver & Roberts 2010). Despite the shared property of projection, CIs differ from other projection phenomena with respect to the information status of their contribution. Presuppositions, for instance, refer to established, or *old* information, whereas CIs contribute *novel* information to the discourse, like at-issue content. Here, we propose a unidimensional analysis of CIs and at-issue content, which highlights the similarity in projection behavior of CIs, presuppositions, and anaphora. This analysis treats CIs as 'piggybacking' on their anchor; they introduce an anaphoric dependency on the interpretation site of their anchor, while at the same time requiring their anchor to refer to a specific referent in the discourse context. CIs are thus elaborations on the description of the referent referred to by their anchor. This analysis of CIs is formalized in Projective Discourse Representation Theory (PDRT; Venhuizen, Bos & Brouwer 2013), a representational framework in which the property of projection is accounted for by explicitly distinguishing between the introduction and interpretation site of semantic content. Our formal analysis explains the interpretational differences between CIs, presuppositions, anaphora, and at-issue content, without stipulating a fundamental distinction between them.

Keywords: conventional implicatures, presuppositions, projection, DRT, PDRT

1 Introduction

The property of *projection* has posed a challenge for semantic theories because of its apparent non-compositional nature. Projection refers to the indifference of semantic content to the syntactic scope of its embedding operators, such as negation, implication, modal operators, and interrogative constructions. Presuppositions have generally been considered the most paradigmatic class of projection phenomena. However, since Potts' redefinition of the class of Conventional Implicatures (CIs; Potts 2003, 2005), interest has shifted toward a broader class of phenomena that

project, including CIs, presuppositions, and anaphora (see, e.g., Simons et al. 2010). Critically, these projection phenomena can be differentiated from *at-issue* content, since they convey backgrounded information. On the other hand, despite their shared property of projection, presuppositions and conventional implicatures also exhibit clear differences in terms of their discourse contribution; while presuppositions signal established, or *old* information (from the speaker's perspective), CIs signal *novel* information, in the sense that by using a CI the speaker communicates to the hearer that he is introducing some additional, backgrounded information that may be new to the hearer. Consider an example of a presupposition in (1a) and a CI in (1b).

(1) a. It is not the case that John and his sister went to the party.
     b. It is not the case that John, who has a sister, went to the party with her.

It follows from both (1a) and (1b) that John has a sister, despite the fact that the constructions that trigger this content syntactically occur within the scope of a negation operator. In (1a) the possessive construction *his sister* triggers a presupposition, i.e., it is taken for granted that John has a sister, while in (1b) the appositive construction signals this same piece of information to be backgrounded but *novel*. The challenge for a formal theory of projection phenomena, now, is to account for the projection behavior of presuppositions and CIs in a unified way, while appreciating their information structural differences. This is not only interesting from a theoretical perspective, but also aids practical approaches to semantic interpretation, such as Information Extraction (see Karttunen & Zaenen 2005).

In this paper, we describe a unified analysis of projection phenomena, formalized in the framework of Projective Discourse Representation Theory (PDRT; Venhuizen et al. 2013). This framework was specifically developed to deal with the projection behavior of presuppositions, without introducing an extra dimension of meaning (as in Geurts & Maier 2003, 2013), or a separate stage of processing (as in van der Sandt 1992). Here, we will show that we can account for CIs within this framework as well, by explicitly incorporating constraints reflecting the information structure of the discourse as part of the semantic representations of PDRT. The analysis of CIs is based on the observations that (i) CIs always attach to an anchor, (ii) this anchor refers to a specific referent in the discourse context, and (iii) CIs project to the same interpretation site as their anchor.

This paper is organized as follows: in section 2, we first discuss the projection behavior of CIs, and propose an analysis in which CI projection is based on discourse anchoring. In section 3, we then describe the semantic framework of PDRT, and show how it can account for the projection behavior of presuppositions. Next, in section 4, the PDRT analysis is extended to account for CIs, on the basis of the analysis presented in the previous sections. The coverage of our novel analysis is

discussed in section 5, and section 6 concludes the paper.

2   The projection behavior of conventional implicatures

CIs were first described by Grice (1975) as referring to implicatures that get their meaning by virtue of the conventional meaning of words. Grice took the expressions "therefore" and "but" to be prime examples of CI triggers, because their usage commits the speaker to the conventional meaning of the expression (i.e. implicating *consequence* and *contrast*, respectively). Potts (2003, 2005) agrees with Grice that some expressions trigger conventional implicatures that commit the speaker to the conventional meaning of the expression, while being logically independent of what is said. However, following Bach (1999), he rejects "but" and "therefore" as CI triggers. Instead, Potts identifies a new class of expressions that trigger CIs, which can be categorized into two main groups: *supplemental expressions*, including appositives, non-restrictive relative clauses (NRRCs) and parenthetical adverbs, and *expressives*, including expressive attributive adjectives, epithets and honorifics. Example (2) illustrates some of these CI triggers (from Potts 2003):

(2) a. Ames, the former spy, is now behind bars. APPOSITIVE
     b. Ames, who stole from the FBI, is now behind bars. NRRC
     c. Ames was, as the press reported, a successful spy. AS-CLAUSE
     d. Fortunately, Beck survived the descent. PARENTHETICAL ADVERB
     e. Frankly (speaking), Ed fled. UTTERANCE MODIFIER
     f. I hate your damn dog! EXPRESSIVE ADVERB
     g. That bastard Conner got promoted. EPITHET
     h. Yamada sensei - ga   o   - warai - ni   nat - ta.
         Yamada teacher - NOM HON - laugh - DAT be - PERF
         'Professor Yamada laughed.' HONORIFIC

In each of the examples above, the speaker is committed to the contribution made by the underlined content; for example, the supplemental in (2b) conveys that Ames stole from the FBI, the expressive in (2f) conveys the speaker's negative attitude toward the addressee's dog, and the honorific in (2h) conveys the social status of Professor Yamada relative to the speaker. Crucially, the CI content contributes novel information that is backgrounded (i.e., not at-issue), and hence projects.

2.1   CIs, presuppositions, and at-issue content

Most formal analyses of CIs have focused on separating their contribution from at-issue content, e.g., by introducing separate meaning dimensions (see, e.g., Potts 2005). Recently, however, converging evidence has emphasized a close interaction

between CIs and other types of content, motivating a *unidimensional* analysis of CIs and at-issue content (see, e.g., Nouwen 2007; Amaral, Roberts & Smith 2007; AnderBois, Brasoveanu & Henderson 2010; Schlenker 2013; Koev 2014). Ander-Bois et al. (2010), for instance, point out that the strong separation between at-issue content and CIs into different meaning dimensions is challenged by the observation that various semantic phenomena "cross the meaning boundary" between these different types of content. In particular, anaphoric dependencies, presuppositions, as well as ellipsis from the at-issue part of a sentence, can be resolved within the CI, as well as the other way around. This is illustrated in (3) (from AnderBois et al. 2010).

(3) $John_x$, who nearly killed a_y woman with his_x car, visited her_y in the hospital.

In this example, the anaphoric expression *his* that occurs within the CI, is bound by the (presuppositional) antecedent *John*, which is introduced as part of the at-issue content. Moreover, the indefinite *a woman*, introduced in the CI, serves as the antecedent for the pronoun *her*, which occurs in the at-issue part of the sentence. This bidirectional dependency motivates a unidimensional and incremental analysis of CIs and at-issue content. Critically, such an analysis should capture the similarities between these types of content, as well as their differences.

    CIs are similar to at-issue content in terms of the contribution they make to the discourse, as illustrated in the following example (from Nouwen 2007).

(4)   a.   Jake, a famous Dutch boxer, lives in Utrecht.
        b.   Jake lives in Utrecht. He is a famous Dutch boxer.

The appositive in (4a) and the second sentence of (4b) make a similar novel contribution to the discourse (that Jake is a famous boxer). However, the difference between the CI contribution and the at-issue content is that the CI introduces backgrounded information. In this respect, CIs are similar to presuppositions, which by definition contribute information that is not at-issue, and therefore backgrounded. Presuppositions differ from CIs, however, in that they can *bind* to (dynamically) accessible antecedents, like anaphora (following van der Sandt 1992). CIs do not allow for such binding, since they signal novel content. This explains why CIs cannot be cancelled, as illustrated in (5) (from Koev 2014).

(5)   a.   If Betty slapped Fred, then she regrets that she slapped him.
        b. ??If Obama is a socialist, then the President, who is a socialist, will raise
              taxes on the rich.

In (5a), the verb *regrets* triggers the factive presupposition that Betty slapped Fred, but since the content of this presupposition occurs as the antecedent of the conditional, the presupposition is cancelled, and therefore the entire sentence does not

presuppose that she slapped him. In contrast, (5b) shows that a similar 'cancellation' of the conventional implicature (i.e., that the President is a socialist) renders the sentence infelicitous. The projected content of the CI matches the content of the conditional statement, just like in (5a), but the CI contribution cannot be bound by this antecedent. This is because the information that the President is a socialist is signalled to be novel by the use of the appositive construction, while this information is already contributed to the discourse by the antecedent of the conditional.

In summary, converging evidence motivates a unidimensional analysis of CIs and at-issue content. This analysis should account for the fact that both types of content contribute *novel* information, contrasting them with presuppositions. On the other hand, the analysis should also account for the fact that CIs, like presuppositions, contribute *backgrounded* information, and hence project. In what follows, we will derive such an analysis, in which CIs project because they 'piggyback' on an anchor. This piggybacking is due to the fact that CIs provide an elaboration on the referent referred to by their anchor.

2.2   CI projection as anchoring

CIs contribute novel information to the discourse that is not part of the at-issue content. As such, CIs place two constraints on the context in which they occur: (i) CIs need to attach their contribution to some part of the discourse, called an *anchor*, and (ii) this anchor must be backgrounded, which means that it must refer to a *specific* referent in the discourse. We motivate both of these requirements below.

(i) CIs attach to an anchor.   It has been observed before that supplemental CIs systematically occur with a syntactic anchor, which allows them to connect their content to the main contribution of the discourse (e.g., Del Gobbo 2003; Nouwen 2007; Heringa 2012). Similarly, an antecedent is required for an expressive CI to connect its subjective content to. This way of discourse anchoring contains a clear element of anaphoric binding. However, CIs make a different contribution to the discourse than anaphora proper; whereas anaphora are used to express identity to their antecedent, CIs contribute novel information about the anchor. In the examples in (6) (from Potts 2003), we have highlighted the anchor of various supplemental and expressive CIs. Note that these examples only include CIs that have a nominal anchor. We will come back to the issue of non-nominal anchors in the discussion in section 5.

(6) a. **Edna**, a fearless leader, started the descent.
     b. **Chuck**, who killed a co-worker, is in prison.
     c. As the judge wrote, **Chuck** agreed that the verdict was fair.
     d. Every Democrat advocating **a proposal for reform** says
         the stupid thing is worthwhile.
     e. **Ame** ga   huri - masa - ta.
         rain NOM fall - HON - PERF
         'It rained.'

In most of these sentences, the anchor can be straightforwardly identified via syntactic attachment, as in (6a) and (6b), or referential dependency, as in (6d) and (6e). Interestingly, however, the anchor in (6c) is not the entire proposition that is subordinated by the *as*-clause, but rather its subject. This is motivated by the observation that we can re-write (6c) as follows:

(6) c'. **Chuck**, about whom the judge wrote that he agreed that the verdict
          was fair, agreed that the verdict was fair.

The *as*-clause in (6c) makes the same contribution as the supplemental clause in (6c'), in which the complement proposition (Chuck agreed that the verdict was fair) is duplicated. Therefore, the CI in (6c) obtains the same anchor as the supplemental clause in (6c'), namely *Chuck*, the subject of the complement proposition.

(ii) The anchor of a CI refers to specific referent. In order to provide novel information that is not at-issue, CIs must attach to an anchor that is backgrounded itself. This means that the anchor must refer to a specific referent that has already been established in the current discourse context. To see why this is the case, consider example (7).

(7) a. John wants to go to a restaurant.
     b. John wants to go to a restaurant, which is in the city center.

On the most intuitive reading of (7a), the indefinite *a restaurant* is interpreted non-specifically (i.e., John wants to go to some restaurant, but does not have a specific one in mind). Now, if an NRRC is attached to this indefinite, as in (7b), the non-specific reading disappears completely (i.e., John now wants to go a specific restaurant, which—by the way—is in the city center). This is because the CI triggered by the NRRC signals novel information that is not at-issue, and therefore *non-restrictive*. Any novel information that is contributed to a non-specific referent will, by contrast, be restrictive and hence at-issue, as it helps in determining the referent within the common ground. Thus, the indefinite in (7b) can only be interpreted specifically,

since the supplemental clause would otherwise contribute restrictive information (as it would restrict the set of restaurants that John wants to go to). In this case, the supplemental clause would not be an NRRC and therefore it would not trigger a CI.

Now, consider the examples in (6) again. The anchors in (6a-c) are all proper names, and trigger presuppositions that project out of the local context, and hence refer to specific referents. Similarly, it follows from (6e) that the existence of rain, as an object for honoring, is presupposed. Interestingly, however, the indefinite description in (6d) is still ambiguous between a specific and a non-specific reading; the sentence can be interpreted both as referring to a particular proposal for reform (a presupposition trigger), or to *any* proposal for reform. In the latter case, the CI triggered by the epithet *the stupid thing* conveys the speaker's general characterization of Democratic proposals for reform as stupid (Potts 2003: 21). Critically, in this case the anchor of the CI is non-specific within the global context, while it is specific with respect to the context in which the CI is introduced; the universal quantifier introduces an embedded context in which for any *specific* proposal for reform it holds that if a Democrat advocates it, he says that the stupid thing is worthwhile. In other words, within the local context of the CI trigger, the object referred to by the anchor is specific, since it is introduced in a different, accessible context (namely, in the antecedent of the conditional triggered by the use of *every*). This highlights an interesting difference between the classes of supplemental and expressive CIs. Supplemental CIs always occur within the same context as their anchors, because they syntactically attach to them. As such, the anchor of a supplemental CI must project in order to be specific with respect to the local context. This is illustrated in example (8), where the existence presupposition triggered by *the king of France* cannot be cancelled in case a CI is attached to the trigger, as in (8b). In contrast, (8a) shows that the bare anchor does allow such cancellation; this is because the presupposition can be accommodated locally (i.e., within the scope of the negation).

(8)   a.   It's not the case that the king of France is bald, because he doesn't exist.
        b.   #It's not the case that the king of France, <u>who lives in Paris,</u> is bald, because he doesn't exist.

Expressive CIs, by contrast, may attach to a non-projecting anchor, as long as it is introduced in a non-local, accessible context. In this case, the referent referred to by the anchor is at least specific with respect to the context in which the CI is introduced, and therefore the CI is felicitous.

Putting the pieces together: CIs are piggybacking on their anchor.   Given that CIs attach to an anchor, which refers to a specific referent in the discourse context, we can explain the projection behavior of CIs. Since the anchor is established within

the current context, the CI serves as an elaboration on the referent referred to by the anchor, and will hence be accessible in the discourse wherever the anchor is accessible. This means that the CI content is contributed directly to the context where the referent of the anchor is introduced or accommodated. The CI thus *inherits* the anchor's accommodation site; in other words, the CI is 'piggybacking' on its anchor. Here, our analysis clearly differs from the one proposed by Nouwen (2007), who treats appositives as introducing an implicit variable that is bound by the anchor. We make the stronger claim that appositives, and other CIs, not only introduce a discourse anaphoric dependency to their anchor, but also inherit its accommodation site; they are *projection-anaphoric*.

Projection-anaphoricity explains the tendency of CIs to project to the global discourse context; presuppositional anchors introducing specific referents are often discourse-new, and thus accommodated at the global discourse context (see Poesio & Vieira 1998). Importantly, projection-anaphoricity also explains why in some cases CIs do *not* project to the global discourse level. This is illustrated in (9) (from Amaral et al. 2007).

(9)   Joan is crazy. She's hallucinating that some geniuses in Silicon Valley have invented a new brain chip that's been installed in her left temporal lobe [...]. Joan believes that **her chip**, which she had installed last month, has a twelve year guarantee.

In this example, the CI triggered by the NRRC does not project to the global discourse context, since it is clear that the speaker does not want to convey that Joan *actually* had a chip installed last month. Note, however, that the direct anchor of the CI is a possessive construction (*her chip*). This construction triggers a presupposition that is bound (in the sense of van der Sandt 1992) to the referent introduced by *a new brain chip*, which is also embedded under an epistemic verb (*hallucinate*). Thus, the anchor is not accommodated at the global discourse level, and therefore the CI that is attached to the anchor is not either. Yet, the fact that the anchor is specific within the local context, allows it to felicitously occur with a (non-restrictive) CI.

In what follows, we will show that the analysis of CIs as piggybacking on their anchor can be formalized using the framework of Projective Discourse Representation Theory (Venhuizen et al. 2013).

3   Toward a formal framework of projection

Despite the prevalence of projection phenomena in discourse, semantic formalisms have mostly treated the property of projection as a deviation from standard meaning construction. Van der Sandt (1992), for instance, proposes an account of presupposition projection in terms of anaphora resolution, which he formalizes in Discourse

Representation Theory (DRT; Kamp 1981; Kamp & Reyle 1993). On this account, projection is carried out by means of a two-stage resolution procedure, in which presuppositions are only resolved after discourse composition has been completed. In order to eliminate such a *post hoc* analysis of projection phenomena, we have recently proposed Projective Discourse Representation Theory (PDRT; Venhuizen et al. 2013), an extension of DRT in which projection is part and parcel of standard meaning construction. PDRT is a wide-coverage semantic formalism that inherits the DRT analysis of a wide range of linguistic phenomena, including anaphora and tense (see Kamp 1981), quantification and plurality (see Kamp & Reyle 1993), attitude reports (see Asher 1986; Zeevat 1996; Maier 2009), and discourse structure (see Asher & Lascarides 2003). This wide applicability is underlined by the adoption of PDRT as the formalism of choice in the Groningen Meaning Bank (GMB), a large-scale corpus of automatically-derived deep semantic representations (Basile, Bos, Evang & Venhuizen 2012; Bos, Basile, Evang, Venhuizen & Bjerva to appear).

3.1 Projective Discourse Representation Theory

PDRT extends classic DRT by making an explicit distinction between the introduction and interpretation site of semantic content through the use of *projection variables*; all basic structures, called Projective Discourse Representation Structures (PDRSs), are associated with a *label*, and all referents and conditions are assigned a *pointer*. In a PDRS, all semantic content occurs 'in situ', thereby reflecting the linguistic surface form, and hence the introduction site of the content. The interpretation site, on the other hand, is determined by the pointers, which can be bound by the label of an accessible PDRS. Critically, the addition of projection pointers allows us to distinguish between projected and at-issue content, without introducing separate meaning dimensions (cf. Layered DRT; Geurts & Maier 2003, 2013).
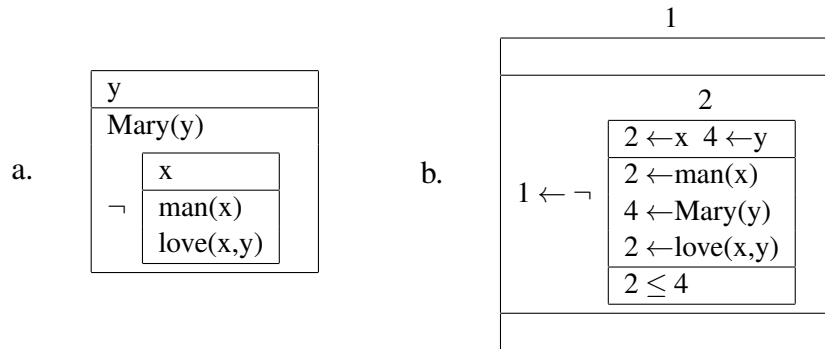
Pointers that are bound by the label of the PDRS in which the content is introduced—the *local* PDRS—indicate at-issue content. Projected content, on the other hand, is either indicated by pointers that are free variables, or by pointers that are bound by the label of some non-local, accessible PDRS. Accessibility is determined in terms of DRS subordination, following Kamp & Reyle (1993): (P)DRS k_1 is *accessible* from (P)DRS k_2 iff k_1 directly or indirectly subordinates k_2. Subordination means that k_2 is part of a condition in k_1, or k_1 serves (directly or indirectly) as the antecedent of k_2. In addition, each PDRS is enriched with a set of 'Minimally Accessible Projection contexts' (*MAPs* for short), which reflect additional constraints on the accessibility of PDRS contexts, in particular those indicated by free pointers (for a similar use of additional accessibility constraints, namely for defining unresolved DRSs, see Reyle 1993, 1995). Together, the classic accessibility constraints and the MAPs create a partial order over PDRS contexts,

which means that each projected context is accessible from at least one (sub-)PDRS that is introduced in the discourse. This assures that the discourse representation remains coherent. Note that the MAPs are an extension to the definition of PDRSs proposed in Venhuizen et al. (2013), which will prove to be a crucial feature for the formalization of the projection behavior of CIs, described in section 4.

3.2 Presuppositions in PDRT

To see how presuppositions are represented in PDRT, consider example (10). This example shows a simple sentence containing a presupposition trigger (the proper name *Mary*), with the corresponding DRS in (10a) and PDRS in (10b).

(10) No man loves Mary.

a.

$$
\begin{array}{|l|}
\hline
y \\
\hline
\text{Mary}(y) \\
\neg \;
\begin{array}{|l|}
\hline
x \\
\hline
\text{man}(x) \\
\text{love}(x,y) \\
\hline
\end{array} \\
\hline
\end{array}
$$

b.

$$
\begin{array}{|c|}
\hline
1 \\
\hline
\\
\hline
1 \leftarrow \neg \;
\begin{array}{|l|}
\hline
2 \\
\hline
2 \leftarrow x \quad 4 \leftarrow y \\
2 \leftarrow \text{man}(x) \\
4 \leftarrow \text{Mary}(y) \\
2 \leftarrow \text{love}(x,y) \\
2 \leq 4 \\
\hline
\end{array} \\
\hline
\\
\hline
\end{array}
$$

As illustrated in this example, PDRSs extend DRSs with projection variables; all PDRS contexts are associated with a label, denoted by the integer on top of the PDRS boxes, and the referents and conditions are all associated with a pointer, denoted by the integer preceding the '←' operator. Finally, the MAPs are shown at the bottom of each PDRS box.

In the PDRS in (10b), all at-issue content is associated with a pointer that is bound by the local PDRS (with label 2), meaning that it should be interpreted within that context. In contrast, the presupposition triggered by the proper name *Mary* (namely, that there exists some person named 'Mary') has a pointer that occurs free, since there is no accessible PDRS that has 4 as its label. This indicates that the presupposition still needs to be resolved to some appropriate context. The only information available in this PDRS is that the accommodation site should be accessible from the local PDRS; this is indicated by the constraint '2 ≤ 4' in the set of MAPs, which means that context 4 is either the same as or higher than local context 2 in the accessibility chain. On top of the semantic constraints provided by the set of MAPs, pragmatic constraints may be employed to resolve the accommodation site

of projected content. For instance, provided the pragmatic assumption that global accommodation is preferred over local accommodation (see van der Sandt 1992; Geurts 1999), we may conclude that the presupposition in (10) will be resolved to the global context (i.e., the PDRS with label 1); this results in an interpretation that is equivalent to the DRS shown in (10a). Despite this interpretational equivalence, it should be noted, however, that the PDRS provides a *richer* representation than the DRS; in contrast to the DRS resulting from van der Sandt's (1992) analysis of presupposition projection, presuppositions remain discernible from at-issue content in the PDRS (namely, on the basis of their pointer).

Moreover, in contrast to the DRS representation, the PDRS representation directly corresponds to the linguistic surface form, as the presupposition trigger *Mary* syntactically occurs within the scope of the negation. This congruity between form and meaning representation facilitates the compositional construction procedure, since no semantic content needs to be moved within the representation, as in van der Sandt's (1992) account. Instead, at-issue and projected content behave compositionally the same (in the sense of Muskens 1996), except for their pointers; during composition, the pointers of at-issue content remain bound by the local context, whereas the pointers of projected content are free variables, which may remain unbound, or become bound later on during discourse construction (for a more elaborate description of this compositional construction process, see Venhuizen et al. 2013; Venhuizen, Bos, Hendriks & Brouwer in prep.).[1]

4 Implementing the projection behavior of CIs

PDRT enriches DRT with an inherent treatment of projection, which has been shown to provide a parsimonious account of presupposition (Venhuizen et al. 2013). In what follows, we will show that PDRT can also account for the projection behavior of conventional implicatures, by implementing the analysis of CIs as piggybacking on their anchor.

4.1 CI anchoring in PDRT

The MAPs of a PDRS impose accessibility constraints on projected contexts. Presuppositions pose minimal constraints on their accommodation site, requiring only that it is accessible from their introduction site. This is formalized using the *weak*

---

1 We have implemented PDRT (and classic DRT) as a Haskell library, called PDRT-SANDBOX. This library incorporates machinery for representing PDRSs and DRSs, translations from PDRT to DRT and first-order logic, composition via different types of merge, and the definition of unresolved structures using Montague Semantics. PDRT-SANDBOX is available at: http://hbrouwer.github.io/pdrt-sandbox/.

subordination constraint ($p_1 \leq p_2$), which means that PDRS context p_2 is either the same as p_1 or resolved to some higher context that is accessible from p_1, thereby licencing local, global and intermediate interpretations of presuppositions (in line with van der Sandt 1992; Geurts 1999). Critically, MAPs can also place stronger constraints on contextual accessibility: *strict subordination* ($p_1 < p_2$) indicates that p_2 is accessible from p_1, but not identical to it, and *identity* ($p_1 = p_2$) indicates that p_1 and p_2 refer to the same projection context. We can use these latter constraints to account for the projection behavior of CIs.

CIs provide a backgrounded elaboration on the description of the referent referred to by their anchor. As such, they *depend* on the interpretation site of this referent. To capture this dependency in terms of a MAP constraint, identity is required between the interpretation site of the CI and that of its anchor. Moreover, as CIs require their anchor to be specific, they *constrain* their own interpretation site, and thereby that of their anchor, to be non-local with respect to their introduction site; this assures that the referent referred to by the anchor is established with respect to the introduction context of the CI. This constraint can be formalized by requiring the interpretation site of the CI to strictly subordinate its introduction site. Example (11), which shows the PDRS for example (4a) above, illustrates this PDRT representation of CIs.[2]

(11)   Jake, a famous boxer, lives in Utrecht.

$$
\begin{array}{|l|}
\hline
\multicolumn{1}{c}{1} \\
\hline
2 \leftarrow x \quad 3 \leftarrow y \\
2 \leftarrow \text{Jake}(x) \\
3 \leftarrow \text{famous\_boxer}(y) \\
3 \leftarrow y=x \\
1 \leftarrow \text{lives\_in\_Utrecht}(x) \\
1 \leq 2 \quad 1 < 3 \quad 3 = 2 \\
\hline
\end{array}
$$

In this example, like in (10), the presupposition (triggered by the proper name *Jake*) is associated with a free pointer, which is related to its introduction context via weak subordination ($1 \leq 2$). The CI triggered by the nominal appositive (*a famous boxer*) is also assigned a free pointer. The pointer of the at-issue content (*lives in Utrecht*), by contrast, is bound by the local context. Critically, the CI also introduces the two aforementioned accessibility constraints. First, the interpretation site of the CI strictly subordinates its introduction context ($1 < 3$), and secondly, the interpretation site of the CI is equated with the interpretation site of its anchor ($3 = 2$). In this way, the CI indirectly constrains the interpretation site of the anchor. Since the CI requires identity to the interpretation site of the anchor *and* projection, the presupposition

---

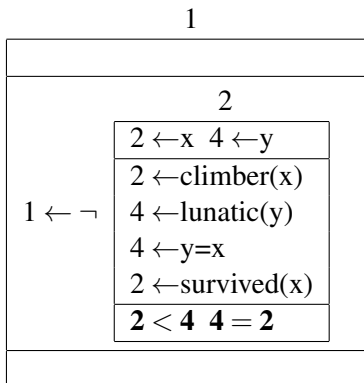2 For reasons of brevity, the CI content and the at-issue content are represented as a single condition.

triggered by the anchor can only be accommodated non-locally (thus, the MAP constraint $1 \leq 2$ is strengthened to $1 < 2$).

It should be noted that the MAP constraints in (11) cannot be satisfied within the current discourse representation. This is the case because the CI requires a non-local context to accommodate to $(1 < 3)$, but no such context is available in this representation. Nonetheless, the PDRS shown in (11) is felicitous. This is because PDRSs (like DRSs) are considered a *partial* representation of the discourse; they will be interpreted with respect to some larger model, in which the unresolved presuppositions can be verified (see Kamp & Reyle 1993; see also the analysis of indexicals proposed by Hunter 2013, where an additional global DRS is introduced to which indexical expressions accommodate).

4.2 Predicting CI infelicity in PDRT

The MAP constraints introduced by CIs generate straightforward predictions about CI (in)felicity. Example (12), for instance, shows that the constraints predict infelicity when a CI occurs with a non-specific anchor (from McCawley 1998).

(12)   #No climber, a lunatic, survived.

1	
$1 \leftarrow \neg$	2

Within box 2:
$2 \leftarrow x$ $4 \leftarrow y$
$2 \leftarrow$ climber(x)
$4 \leftarrow$ lunatic(y)
$4 \leftarrow y=x$
$2 \leftarrow$ survived(x)
$2 < 4$ $4 = 2$

In this example, the non-specific anchor is locally accommodated. As a consequence, its interpretation site is the same as the local context of the CI. This leads to a contradiction within the accessibility constraints; the constraints introduced by the CI simultaneously require the CI to project out of its local context (strict subordination; $2 < 4$), and its interpretation site to be equal to that of its anchor (identity; $4 = 2$). These constraints ($2 < 4$ and $4 = 2$) cannot be simultaneously satisfied, and therefore the PDRS is infelicitous.

The PDRT analysis also predicts that CIs cannot be cancelled. This is illustrated in (13), which shows the PDRS for example (5b) from above.

(13) ??If Obama is a socialist, then the President, <u>who is a socialist,</u> will raise taxes.
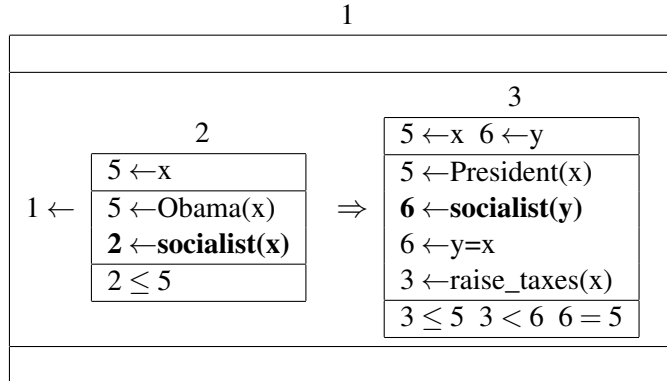
In this example, the presupposition triggers *Obama* and *the President*, introduced in respectively the antecedent and the consequent of an implication, refer to the same entity. This is reflected by the fact that these conditions affect the same referent and are assigned the same pointer.[3] As such, the contribution made by the CI (that the president is a socialist) is the same as the contribution in the antecedent of the implication (that Obama is a socialist). This equivalence renders the antecedent of the implication void (i.e., not locally informative; cf. van der Sandt 1992). This is because the identity constraint ($6 = 5$) introduced by the CI forces the CI content to project out of the consequent to the interpretation site of its anchor. As a result of this piggybacking, the content contributed by the antecedent is already established within the context in which it is introduced; this is because the MAP constraint in the antecedent ($2 \leq 5$) indicates that the interpretation site of the CI content is accessible from the antecedent's local context. Hence, because CIs piggyback on their anchor, they cannot be cancelled unless their anchor is also cancelled. In (13), the latter is not the case, and the example is therefore infelicitous.

4.3 Expressive CIs and speaker-orientedness

In Potts' (2005) redefinition of the class of CIs, a central property of CIs is *speaker-orientedness*. The importance of this property is particularly emphasized in the analysis of expressive CIs, which contribute an opinion (the speaker's, in most cases) to the at-issue content of a sentence. In PDRT, the attribution of subjective attitudes

---

3 In PDRT, uniqueness of referents is determined on the basis of their pointer, whereas in DRT it is determined on the basis of their introduction site. Therefore, referents may be introduced more than once in PDRT, which is not allowed in DRT due to ambiguous binding. See Venhuizen et al. (in prep.) for details.

is a pragmatic process, driven by context; both the possible belief-bearers and the polarity of their beliefs are derived by pragmatic inferencing and the application of world knowledge. This is exemplified by (14) (adapted from Kratzer 1999), where the choice of verb affects to whom the subjective content of the epithet (*that bastard*) is contributed.

(14)   a.   Father said he would not allow me to *hit* that bastard Webster.
        b.   Father said he would not allow me to *marry* that bastard Webster.
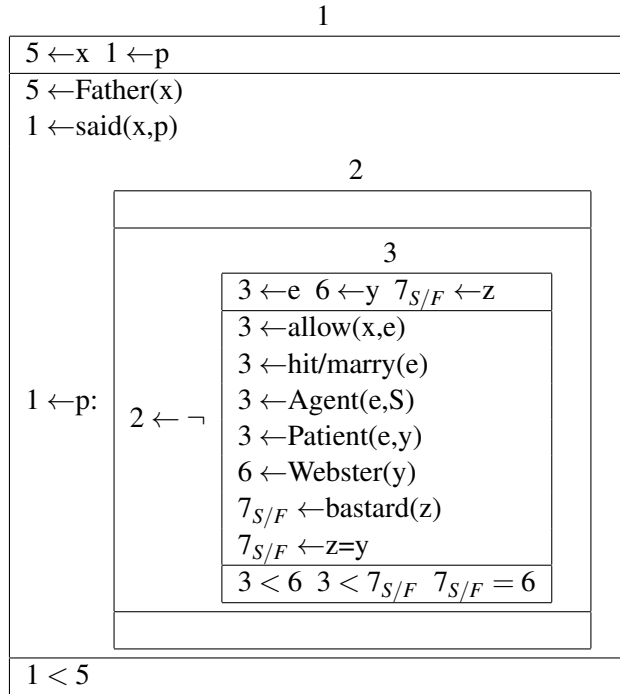
In (14a), the verb *hit* suggests a negative attitude of the speaker toward the person named 'Webster', and therefore world knowledge will tell us that the speaker is the most likely bearer of the (negative) subjective content triggered by the CI (*that bastard*). In contrast, this same subjective content will most likely be attributed to the father in (14b), since the speaker's attitude toward Webster seems to be positive; the context created by the verb *marry* suggests that the speaker wants to marry Webster.

Pragmatics-driven subjectiveness can be implemented in PDRT by exploiting the unresolved nature of free pointers in the model-theoretic interpretation. In the model-theoretic interpretation of classic DRT (see, e.g., Kamp 1981; Kamp & Reyle 1993; Kamp, van Genabith & Reyle 2011), a DRS is considered a *partial model*, representing the information conveyed by an utterance. In order to determine the truth-conditional interpretation of this DRS, the partial model needs to be embedded in a *total model*. This embedding is done via an *embedding function* that maps the set of discourse referents in the universe of a DRS K, onto elements in the domain of a total model M. This mapping is done in such a way that all conditions of K are verified in M (with respect to some world w). A model-theoretic interpretation of PDRT can be constructed in a similar way. However, the embedding function for PDRSs needs to have proper machinery for dealing with projection variables. Content associated with bound projection variables should be interpreted in its binding context. Content associated with free variables, on the other hand, still has an undetermined interpretation site, and should be resolved to an appropriate context, on the basis of the semantic constraints in the MAPs, and potential additional pragmatic constraints. One possible interpretation site for unresolved content is the global discourse context, which reflects the common ground of the speaker and hearer. Alternatively, unresolved content may be interpreted in a *subjective* discourse context, which provides a means of dealing with speaker-orientedness in PDRT.

Speaker-orientation is formalized in PDRT by assigning each discourse agent (e.g., the speaker, the hearer, and possible other individuals introduced in the discourse) a subjective model M_s, which describes the current state of affairs in the world, as believed by that agent. Critically, this subjective model may or may not coincide with the *actual* state of affairs in the world, which is represented in the truth

model M_T. Pragmatic constraints determine whether subjective content associated with a free pointer should be verified with respect to a discourse agent's subjective model M_s, rather than with respect to the truth model M_T. If this is the case, verification of the entire utterance with respect to the truth model M_T entails that this subjective content *only* needs to be verified with respect to the subjective model M_s. Provided this analysis, we can represent the subjective sentence from example (14) as the PDRS shown in (15).

(15)   Father$_F$ said he would not allow me$_S$ to hit/marry <u>that bastard</u> Webster.

<div align="center">

1

$5 \leftarrow x$ $1 \leftarrow p$
$5 \leftarrow$ Father(x)
$1 \leftarrow$ said(x,p)

2

		3	
		$3 \leftarrow e$ $6 \leftarrow y$ $7_{S/F} \leftarrow z$	
		$3 \leftarrow$ allow(x,e)	
		$3 \leftarrow$ hit/marry(e)	
$1 \leftarrow p$:	$2 \leftarrow \neg$	$3 \leftarrow$ Agent(e,S)	
		$3 \leftarrow$ Patient(e,y)	
		$6 \leftarrow$ Webster(y)	
		$7_{S/F} \leftarrow$ bastard(z)	
		$7_{S/F} \leftarrow z = y$	
		$3 < 6$ $3 < 7_{S/F}$ $7_{S/F} = 6$	

$1 < 5$

</div>

In this example, there are two relevant discourse agents that may express their subjective beliefs: the speaker and the father, which are associated with subjective models M_S and M_F, respectively. To verify the truth of the entire utterance with respect to a truth model M_T, the subjective content expressed by the CI (*that bastard*), which is associated with a free pointer, should be verified with respect to either M_S (in case of the verb *hit*) or M_F (in case of *marry*). In (15), the subjective model with respect to which the subjective CI content is interpreted, is indicated by means of the subscript (S or F) attached to the free pointer. These 'subjective' projection variables are tuples of pointers and subjective models (e.g., $\langle 7, S \rangle$ or $\langle 7, F \rangle$), that inform the embedding function with respect to which model the subjective content should be interpreted.

In summary, speaker-orientedness is dealt with at the pragmatic, rather than the semantic level in PDRT. This explains the context-dependency of the phenomenon, both on the linguistic context, as illustrated in (14), and on the general discourse context, which includes non-linguistic factors, such as the type of *speaker* (cynical, authoritative, etc.), and the *situation* in which a sentence is uttered (for example, a formal setting or an informal conversation) (see Hoeks & Brouwer 2014). By accounting for CIs within the semantic representation, while attributing their speaker-orientedness to the pragmatic, interpretational domain, the PDRT analysis of CIs intuitively captures the multi-dimensionality of their contribution. This shows that PDRT provides a robust and flexible framework for the representation and interpretation of different types of linguistic phenomena and their associated levels of information.

5 Discussion

Formal analyses of CIs have often focused on differentiating their contribution from at-issue content, as well as from that of other projection phenomena, such as presuppositions. In this paper, we have proposed an analysis of the projection behavior of CIs, which highlights their correspondence to presuppositions and anaphora, and treats them within the same dimension as at-issue content. In this analysis, formalized in the framework of Projective Discourse Representation Theory, CIs piggyback on a specific anchor because their contribution provides an elaboration on the description of the referent referred to by this anchor.

Concerning Potts' (2005) class of CIs, we have shown that our analysis accounts for the (in)felicity of supplemental and expressive CIs that have a nominal anchor. However, Potts also discusses CIs that do not have a nominal anchor, such as utterance modifiers (e.g., "frankly") and expressive adverbs (e.g., "fucking"). If these expressions do indeed trigger CIs, we might have to extend our analysis to non-nominal anchors. The CI status of these expression has, however, been a subject of debate. Amaral et al. (2007: 725-729), for instance, provide an extensive discussion about utterance level modifiers, like "frankly", arguing that they behave differently from other CIs, both on the theoretical and implementational level; they seem to target the *act* of uttering a statement, rather than the utterance itself. Similarly, Geurts (2007) argues that expressive adverbs, like "fucking", have a semantic dimension of their own, thereby differentiating them from, in particular, supplemental CIs. Hence, the analyses of Amaral et al. (2007) and Geurts (2007) cast doubt on the CI status of utterance modifiers and expressive adverbs. Based on the analysis of CIs as piggybacking on their anchor, we may have to draw a similar conclusion; it remains to be seen whether there indeed exist CIs that have a non-nominal anchor.

A core assumption of our analysis of CIs is that they elaborate the description

of a specific referent in the discourse context, as referred to by an anchor. As such, CIs piggyback on their anchor by projecting to its interpretation site. Schlenker (2013), however, argues that there exist CIs that are interpreted *in situ*, and thus do not project to the interpretation site of their anchor. One example, in which he argues this to be the case, is shown in (16).

(16)   [Context: someone made a big mistake at the Department.]
         If tomorrow I called the Chair, who in turn called the Dean, then we would
         be in deep trouble.

According to Schlenker (2013), the supplemental construction *who in turn called the Dean* is an NRRC that triggers a CI, which is interpreted locally, as indicated by the licensing of the past tense in the relative clause. However, as the relative clause introduces an additional conditional statement, rather than an elaboration on the description of the referent referred to by the anchor *the Chair*, we disagree with Schlenker (2013) on the CI status of this contribution.

Nouwen (2014) also argues for the existence of CIs that have an *in situ* interpretation, which he exemplifies by means of (17).

(17)   If a professor, a famous one, writes a book, he will make a lot of money.

According to Nouwen (2014), the indefinite anchor *a professor* can be interpreted non-specifically, while at the same time occurring felicitously with the appositive construction *a famous one*. This means that the CI triggered by the appositive is interpreted *in situ*, namely in the same context as its non-specific, and therefore non-projecting, anchor. However, by virtue of being interpreted locally, the appositive becomes restrictive. As such, it does not provide a backgrounded elaboration on its anchor, and is therefore not a CI.

6   Conclusion

We have proposed an analysis of how and why conventional implicatures project; CIs piggyback on their anchor, which refers to a specific referent in the discourse context. This piggybacking is due to the fact that CIs signal novel, backgrounded information that elaborates the description of the referent referred to by the anchor. We have formalized our analysis in the framework of Projective Discourse Representation Theory. This PDRT analysis provides a unified account of the projection behavior of CIs, presuppositions, and anaphora, and treats their contribution within the same dimension as that of at-issue content.

References

Amaral, Patricia, Craige Roberts & E. Allyn Smith. 2007. Review of 'The Logic of Conventional Implicatures' by Chris Potts. *Linguistics and Philosophy* 30(6). 707–749.

AnderBois, Scott, Adrian Brasoveanu & Robert Henderson. 2010. Crossing the appositive/at-issue meaning boundary. In Nan Li & David Lutz (eds.), *Semantics and Linguistic Theory (SALT) 20*, 328–346. CLC Publications.

Asher, Nicholas. 1986. Belief in discourse representation theory. *Journal of Philosophical Logic* 15(2). 127–189.

Asher, Nicholas & Alex Lascarides. 2003. *Logics of Conversation*. Cambridge University Press.

Bach, Kent. 1999. The myth of conventional implicature. *Linguistics and Philosophy* 22(4). 367–421.

Basile, Valerio, Johan Bos, Kilian Evang & Noortje J. Venhuizen. 2012. Developing a large semantically annotated corpus. In Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, Mehmet Uğur Doğan, Bente Maegaard, Joseph Mariani, Asuncion Moreno, Jan Odijk & Stelios Piperidis (eds.), *International Conference on Language Resources and Evaluation (LREC) 8*, 3196–3200. Istanbul, Turkey: European Language Resources Association (ELRA).

Bos, Johan, Valerio Basile, Kilian Evang, Noortje J. Venhuizen & Johannes Bjerva. to appear. The Groningen Meaning Bank. In Nancy Ide & James Pustejovsky (eds.), *Handbook of Linguistic Annotation – Part Two: Case studies* Text, Speech and Language Technology, Springer.

Del Gobbo, Francesca. 2003. *Appositives at the Interface*: University of California, Irvine PhD Thesis.

Geurts, Bart. 1999. *Presuppositions and pronouns*. Elsevier.

Geurts, Bart. 2007. Really fucking brilliant. *Theoretical Linguistics* 33. 209–271.

Geurts, Bart & Emar Maier. 2003. Layered DRT. Ms.

Geurts, Bart & Emar Maier. 2013. Layered Discourse Representation Theory. In Alessandro Capone, Franco Lo Piparo & Marco Carapezza (eds.), *Perspectives on Linguistic Pragmatics*, 311–327. Springer International Publishing.

Grice, H. Paul. 1975. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Syntax and semantics*, vol. 3 Speech acts, 41–58. San Diego, CA: Academic Press.

Heringa, Herman. 2012. *Appositional Constructions*: University of Groningen PhD Thesis.

Hoeks, John C. J. & Harm Brouwer. 2014. Electrophysiological research on conversation and discourse processing. In Thomas M. Holtgraves (ed.), *The Oxford Handbook of Language and Social Psychology*, chap. 23, 365–386. New York:

Oxford University Press.

Hunter, Julie. 2013. Presuppositional indexicals. *Journal of Semantics* 30(3). 381–421.

Kamp, Hans. 1981. A theory of truth and semantic representation. In Jeroen A. G. Groenendijk, Theo M. V. Janssen & Martin B. J. Stokhof (eds.), *Formal Methods in the Study of Language* 135, 277–322. Mathematisch Centrum.

Kamp, Hans, Josef van Genabith & Uwe Reyle. 2011. Discourse Representation Theory. In Dov M. Gabbay & Franz Guenthner (eds.), *Handbook of Philosophical Logic*, vol. 15, 125–394. Springer Netherlands.

Kamp, Hans & Uwe Reyle. 1993. *From discourse to logic: Introduction to modeltheoretic semantics of natural language, formal logic and Discourse Representation Theory*. Dordrecht: Kluwer.

Karttunen, Lauri & Annie Zaenen. 2005. Veridicity. In Graham Katz, James Pustejovsky & Frank Schilder (eds.), *Annotating, Extracting and Reasoning about Time and Events* (Dagstuhl Seminar Proceedings 05151), 1–9. Dagstuhl, Germany: Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI).

Koev, Todor. 2014. Two puzzles about appositives: Projection and perspective shift. In Urtzi Etxeberria, Anamaria Fălăuş, Aritz Irurtzun & Bryan Leferman (eds.), *Sinn und Bedeutung 18*, 217–234. Bayonne and Vitoria-Gasteiz: University of the Basque Country (UPV/EHU).

Kratzer, Angelika. 1999. Beyond 'Ouch' and 'Oops'. How descriptive and expressive meaning interact. Cornell Conference on Theories of Context Dependency. Comments on a paper by Kaplan.

Maier, Emar. 2009. Presupposing acquaintance: a unified semantics for de dicto, de re and de se belief reports. *Linguistics and Philosophy* 32(5). 429–474.

McCawley, James D. 1998. *The syntactic phenomena of English*. University of Chicago Press.

Muskens, Reinhard. 1996. Combining Montague semantics and discourse representation. *Linguistics and Philosophy* 19(2). 143–186.

Nouwen, Rick. 2007. On appositives and dynamic binding. *Research on Language & Computation* 5. 87–102.

Nouwen, Rick. 2014. A note on the projection of appositives. In Eric McCready, Katsuhiko Yabushita & Kei Yoshimoto (eds.), *Formal approaches to semantics and pragmatics: Japanese and beyond*, vol. 95 Studies in Linguistics and Philosophy, Springer.

Poesio, Massimo & Renata Vieira. 1998. A corpus-based investigation of definite description use. *Computational Linguistics* 24(2). 183–216.

Potts, Christopher. 2003. *The Logic of Conventional Implicatures*: University of California PhD Thesis.

Potts, Christopher. 2005. *The Logic of Conventional Implicatures*. Oxford University Press, USA.

Reyle, Uwe. 1993. Dealing with ambiguities by underspecification. *Journal of Semantics* 10(2). 123–179.

Reyle, Uwe. 1995. On reasoning with ambiguities. In Steven P. Abney & Erhard W. Hinrichs (eds.), *European chapter of the Association for Computational Linguistics (EACL) 7*, 1–8. University College Dublin, Belfield, Dublin, Ireland. The Association for Computational Linguistics.

van der Sandt, Rob. 1992. Presupposition projection as anaphora resolution. *Journal of Semantics* 9. 333–377.

Schlenker, Philippe. 2013. Supplements without bidimensionalism. Ms.

Simons, Mandy, Judith Tonhauser, David Beaver & Craige Roberts. 2010. What projects and why. In Nan Li & David Lutz (eds.), *Semantics and Linguistic Theory (SALT) 20*, 309–327. CLC Publications.

Venhuizen, Noortje J., Johan Bos & Harm Brouwer. 2013. Parsimonious semantic representations with projection pointers. In Katrin Erk & Alexander Koller (eds.), *International Conference on Computational Semantics (IWCS) 10 – Long Papers*, 252–263. Potsdam, Germany: Association for Computational Linguistics. http://www.aclweb.org/anthology/W13-0122.

Venhuizen, Noortje J., Johan Bos, Petra Hendriks & Harm Brouwer. in prep. Harnessing projection: implementing Projective Discourse Representation Theory.

Zeevat, Henk W. 1996. A neoclassical analysis of belief sentences. In Paul J. E. Dekker & Martin J. B. Stokhof (eds.), *Amsterdam Colloquium 10*, 723–742. ILLC, Universiteit van Amsterdam.

Noortje J. Venhuizen
CLCG, University of Groningen
Oude Kijk in 't Jatstraat 26
9712 EK Groningen
The Netherlands
n.j.venhuizen@rug.nl

Petra Hendriks
CLCG, University of Groningen
Oude Kijk in 't Jatstraat 26
9712 EK Groningen
The Netherlands
p.hendriks@rug.nl

Johan Bos
CLCG, University of Groningen
Oude Kijk in 't Jatstraat 26
9712 EK Groningen
The Netherlands
johan.bos@rug.nl

Harm Brouwer
Department of Computational
Linguistics and Phonetics
Saarland University
66123 Saarbrücken
Germany
brouwer@coli.uni-saarland.de