

Entailed vs. Non-Entailed Presuppositions - An Experimental Assessment

Jérémy Zehr & Florian Schwarz

University of Pennsylvania

1. Introduction

There's a current debate about presuppositions about whether different types of triggers need to be distinguished, with various proposals for how to flesh out potential distinctions, e.g., lexical vs. resolution triggers (Zeevat 1992) or soft vs. hard triggers (Abusch 2010) (for a broader and more fine-grained classification, see Tonhauser et al. 2013). Building on Sudo (2012), we compare triggers that he argues to differ in whether their presupposition is also part of what is entailed (e.g., *continue*, *stop*) or not (e.g., *also*), as reflected in their behavior under non-monotonic quantifiers. While the results are largely consistent with Sudo's proposal at first sight, the effects for the alleged entailed presuppositions are not as uniform as expected, and thus may ultimately be more amenable to an alternative explanation based on local accommodation.

1.1 Are presuppositional and entailed content fully distinct or can they overlap?

Presuppositional sentences are commonly thought to contribute meaning on two different levels. For instance, (1) conveys both (1a) and (1b):

- (1) The temperature has continued to increase.
 - a. The temperature was previously increasing
 - b. The temperature is currently increasing

The key motivation for treating these two contributions to the overall meaning separately is that they behave differently under operators like negation: (1a) is not affected by embedding under negation ((2) still conveys (1a)), but (1b) is ((2) no longer conveys (1b)).

- (2) The temperature has not continued to increase.

In representing the meaning of presuppositional sentences, these two contributions are therefore commonly separated as outlined above: the presupposition (1a) is identified as the part that is unaffected by negation (as presuppositions *project*) and the remainder of the meaning is taken to be the ‘regular’, or *entailed*, content that negation applies to (1b). This makes for a straightforward description of what content is targeted by negation, in line with (2) conveying that (1b) is false. But note that the data can also be captured if we assume that the entailed content that negation targets is actually the *conjunction* of (1a) and (1b), as long as (1a) simultaneously retains presuppositional status: taken together, the negated conjunction (3a) and the (projecting) presupposition (1a) entail the negation of (1b).

- (3) a. NOT [the temperature was previously increasing and it is increasing now]
 b. (3a) & (1a) → the temperature is not increasing

In principle, then, both options for characterizing the entailed content – (1b) alone vs. the conjunction of (1a) and (1b) – seem viable as things stand. What might lead us to favor one over the other? On purely conceptual grounds, the former option might seem preferable, given that it separates the two contributions more neatly, and directly captures what is negated without having to consider what is presupposed.

However, various attempts to capture differences between triggers (e.g., Simons 2001, Abusch 2010, Romoli 2014) rely on the second approach, as they are based on the idea that the presuppositional nature of the presuppositions of soft triggers is pragmatically derived. While details vary across accounts, the key shared assumption is that what is traditionally seen as a presupposition is actually introduced as entailed content at the lexical level, and only is accorded presuppositional status via pragmatic reasoning. It is hard to see how any account in the general spirit of these proposals could get off the ground without assuming that what winds up as being presupposed is also part of the entailed content.

Aside from such theoretical motivations, what empirical evidence could bear on the issue? As Sudo (2012) spells out, truth-value judgments would seem to be of little help for most presuppositional sentences in this regard. In most environments, the contribution of an expression at the presuppositional level would not be distinguishable from its contribution at the non-presuppositional level. A presupposition must hold in any scenario where the sentence in question is felicitous and true, consistent with either option. And any scenario in which the presupposition does not hold will yield the sentence infelicitous on either option. As Sudo lays out, this issue arises for presuppositions in monotonic environments in general. However, in non-monotonic environments, e.g. the scope of *exactly one*, the predictions of the two options come apart. Based on this diagnostic, which we lay out in detail in the next section, Sudo (2012) argues that triggers can differ precisely with regards to whether or not their presupposition also features in the non-presuppositional, entailed content.¹ In particular, Sudo (2012) argues gender features on pronouns (the main focus of his work) to be presupposed but not part of the entailed content (for a recent extension

¹Note that similar ideas have been raised in different contexts by Gajewski & I-ta Chris Hsieh (2014), Gajewski (2016) (the latter was in circulation as early as 2011), who argued for a parallel distinction between singular and plural definites to account for their differences concerning the licensing of NPIs. Also see Chierchia (2015) for further development of the role of such contrasts between triggers for NPI-licensing.

to gender features on Greek nouns, see Sudo & Spathas 2015). In contrast, certain other triggers exhibit interpretations that suggest that their presupposition is present at the level of entailed content as well. The main empirical diagnostic for this property that he proposes, which also forms the core of our experiments, is introduced in detail in the next section.

1.2 The *Exactly One* Test

Sudo's main empirical argument that the presupposition introduced by the gender features of pronouns does not contribute to the entailed content is based on their effect under non-monotonic quantifiers such as *exactly one*, as in (4) (Sudo 2012).

- (4) Exactly one student criticized herself (... namely Mary).
- a. Exactly one student is a female who was self-critical.
 - b. Exactly one student was self-critical.

If the gender feature of *her* was entailed in (4), then it should factor into the assessment of *exactly one*, as paraphrased in (4a). But, Sudo argues, this would predict (4) to be compatible with there being a male student that also is self-critical, contrary to intuitions. The intuitively correct reading can be derived if *exactly one* applies to entailing content to the exclusion of the gender feature presupposition, as illustrated in (4b). Moving beyond gender features, however, Sudo argues that other triggers render different results. For example, *stop* seems to favor an entailing reading (5a) over a non-entailing reading (5b): intuitively, the truth of (5) *is* compatible with there being other students that also are not using Mac now, as long as they didn't previously use Mac in the first place.

- (5) Exactly one student stopped using Mac.
- a. Exactly one student used Mac and does not use Mac now.
 - b. Exactly one student does not Mac now.

With regards to other triggers, Sudo tentatively proposes that *even* and additive/hard presupposition triggers in general pattern with gender features and don't entail their presupposition, in contrast to *stop*. The experiments reported here adapt the *exactly one* test to a covered-box design to investigate whether *presuppositional expressions* *i. entail their presuppositions, ii. do not entail their presuppositions, or iii. vary precisely in this respect.*

2. Experiments

2.1 Methods and Design

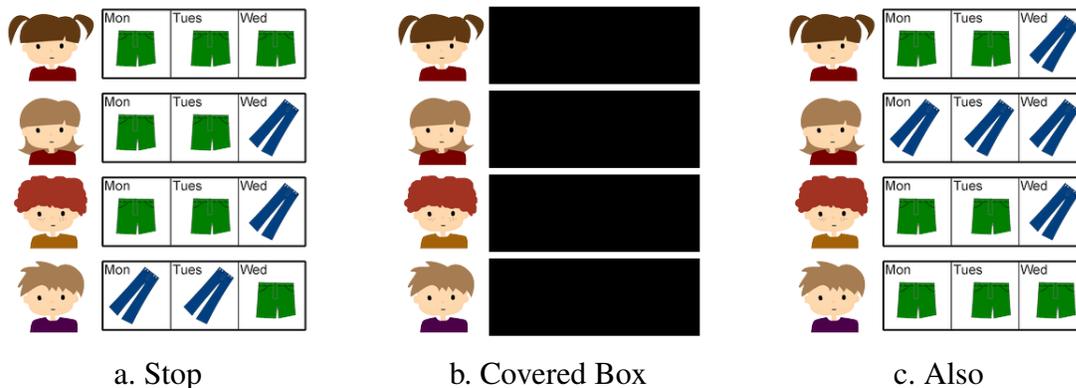
The experiments utilize a picture-matching task with a covered box representing 'hidden' information (Huang et al. 2013, Romoli & Schwarz 2015) and auditory presentation of linguistic stimuli. The images depicted individuals paired with a calendar strip containing iconic representations of events on a given day. The overall task was framed by telling par-

ticipants that groups of four people at a fashion camp had to come up with two alternative plans for wearing clothes. Participants then had the job to determine, based on the reported sentences, which of the two plans the group wound up carrying out. The fully visible plan was to be chosen if it was found to match the linguistic description, and the covered box, with detailed information hidden from view, otherwise. To highlight information on what happened before Wednesday, the Wednesday boxes were not visible during the playback of the context sentence (6a) and only appeared when the test sentence (6b) started to play.

Critical *stop*-trials used pictures such as the one in (7a) along with the covered box (7b), and the sentence in (6b), after an introductory context sentence (6a). (Note that, following Sudo (2012) and in line with results from Chemla (2009), we take (6b) to come with an existential presupposition. This assumption is supported by our results).

- (6) a. At the beginning of the week, some kids wore pants and some wore shorts.
 b. Exactly one kid stopped wearing pants on Wednesday.

(7) Critical Images:



The *exactly one* context allows us to assess whether the presuppositional content also features at the level of entailment: (7a) is consistent with the (existential) presupposition formulated in (8a), as *at least one kid* (indeed, exactly one) wore pants before Wednesday. (8b) correctly describes the situation shown: only one kid wore pants before but not on Wednesday (the bottom-most one). In contrast, (8c) is not compatible with the situation, as *two* kids wear shorts, i.e. not pants, on Wednesday (the top-most and bottom-most ones). Thus, choices of the visible picture indicate an entailing interpretation (8b), whereas covered box choices indicate a non-entailing interpretation (8c).²

- (8) a. At least one kid wore pants before Wednesday.
 b. Exactly one kid wore pants before Wednesday but not on Wednesday.
 c. Exactly one kid didn't wear pants on Wednesday.

²The non-presuppositional contents that we discuss in this paper are not stated in Sudo's terms, as he uses a bi-dimensional representation augmented with a cross-sentential anaphoric device to be able to refer to the same 'at-least-one' individual in the presupposition as in the 'exactly-one' non-presuppositional content.

Entailed vs. Non-Entailed Presuppositions

The second trigger we investigated was *also* (9), paired with (7c) as the critical target image. Note that recordings systematically exhibited stress on *Wednesday* to ensure appropriate association of focus with *also*. Furthermore, the context sentence (6a) ensured the felicity of *also* with regards to a sufficiently salient antecedent event. The logic of the mapping of the candidate interpretations to the critical visible picture was as before, with (9b) representing an interpretation where the presupposition is also part of the entailed content (which is true of (7c)), and (9c) representing the alternative reading where the entailed content does not include the presupposition (which is false).

- (9) Exactly one kid also wore pants [on Wednesday]_F.
- a. At least one kid wore pants before Wednesday.
 - b. Exactly one kid wore pants before and on Wednesday.
 - c. Exactly one kid wore pants on Wednesday.

True and False Control condition images made both versions of the entailed content true and false respectively. Additional Filler conditions were added to balance the rate of expected visible vs. covered choices. In an additional filler condition, the visible picture was only compatible with a potentially available reading where *also* takes scope over *exactly one* (*On Wednesday, too, exactly one kid wore pants*).³ In the visual stimuli, characters were always grouped by gender, with vertical order counterbalanced. In the linguistic stimuli, target sentence reference to the first vs. second clothing item in the context sentence was also counterbalanced, as was the role of male vs. female characters for the sentence's evaluation. Participants saw 6 items per critical condition (for a total of 36, including the *her*-conditions), and 48 filler items, with item-condition pairings counterbalanced across participants in a latin square design. The experiment began with a practice trial.

Two versions of the experiment were run: a first version, where *stop* and *also* (as well as *her*) items were randomly intermixed yielded mixed results for *stop*. A subsequent experiment presented *stop*, on the one hand, and *also* and *her*, on the other hand, in separate blocks to test whether intermixing *stop* with the other triggers was responsible for this. 48 and 36 participants were recruited via Mechanical Turk for the mixed vs. block design respectively. The experiment was implemented via Ibex. Response choices and reaction times (RTs; starting from the playback of the test sentence) served as dependent variables.

2.2 Results

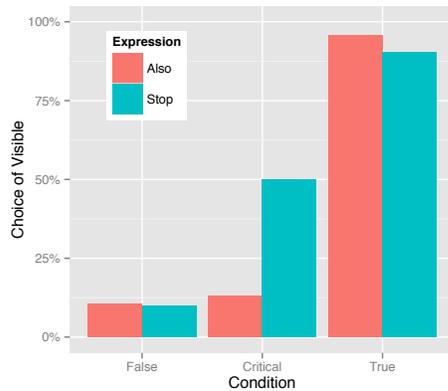
2.2.1 Mixed Design

Participants' accuracy was assessed in the *stop* and *also* fillers and controls. The results in (10a) are based on the 33 participants who had an accuracy of at least 75% on the *also* and *stop* controls and fillers (the overall pattern did not change when the exclusion criterion was

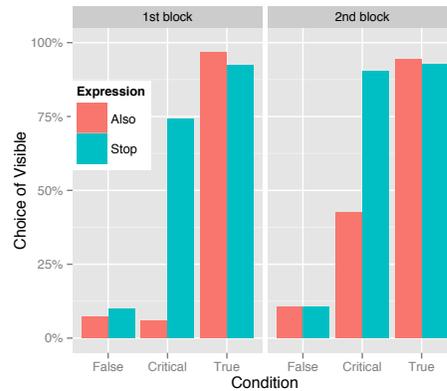
³We also included items testing the gender feature of *her*, given the central role of these for Sudo's proposal. We do not discuss these here, as the results were somewhat mixed, most likely due to insufficient salience of *her* in the auditory stimuli, which easily could be misheard as *a* in many cases.

relaxed). We used mixed-effect logistic regressions models (in the *lme4* package in R) to analyze response patterns by condition, with the maximal random effect structure for items and participants that would converge. *Also-critical* and *also-false* did not differ significantly ($p = 0.4233$) but *stop-critical* did from both *stop-false* and *stop-true* (p 's < 0.01).

(10) *Response rates*



a. *Mixed Design*



b. *Block Design*

To analyze RTs, we ran linear mixed-effect models, predicting the log-transformed RTs based on condition, again with the maximal random effects structure that would converge. Significance was determined via model comparisons.

(11) *Mean RT (in ms) for covered (in black) and visible choices (in gray).*

	<i>False</i>	<i>Critical</i>		<i>True</i>
Also	4881	4899	NA	4309
	<i>Wide</i>			
	4453	4525		
Stop	5436	3939	6002	5433

In the *stop-critical* conditions, where we find a mixed response pattern, visible picture choices were significantly slower than covered box choices ($p < .05$). *Covered* box choices in the *also-critical* and the *also-false* conditions did not differ significantly ($p = 0.57$), but were significantly faster in the *stop-critical* condition than in the *stop-false* condition. Visible picture choices in the *stop-critical* conditions did not differ significantly from those in the *stop-true* condition ($p = 0.52$). To shed further light on the contrast in the *stop-critical* condition, a further comparison with the *also-wide-scope* condition was conducted, which was the only other condition that systematically yielded both *visible* and *covered* choices. An interaction analysis by condition and response revealed a significant interaction, i.e. the contrast was significantly greater in *stop-critical* than in *also-wide-scope*.

2.2.2 Block Design

Results in (10b) reflect the 25 participants with accuracy rates of at least 75% on the *also* and *stop* controls and fillers. Mixed-effect models revealed no significant difference between *stop-critical* and *stop-true* in either block ($p = 0.7$ and $p = 0.8$), but there was a marginal interaction between condition and block ($p = 0.099$). There was no significant difference between *also-critical* and *also-false* in the first block ($p = 0.15$). In contrast, there were significantly more visible picture choices for *also-critical* in the second block ($p < .01$). Looking at both blocks together furthermore revealed a significant interaction between *Condition* and block order for *also* (critical vs. false).

(12) Mean RTs (in ms) for covered (in black) and visible choices (in gray).

	1 st Block			2 nd Block			
	<i>False</i>	<i>Critical</i>	<i>True</i>	<i>False</i>	<i>Critical</i>	<i>True</i>	
Also	5392	4242	(4538)	3809	3967	6376	3691
Stop	(4780)	6115	4966	(4827)	5273		5173

RTs for the *covered* choices in the *also-critical* and in the *also-false* conditions did not significantly differ in either block ($p = 0.45$ and $p = 0.28$). In contrast, RTs for the *visible* choices in the *stop-critical* condition were significantly slower than in the *stop-true* condition in the first block ($p < .05$), but not in the second ($p = 0.59$). In the second block, RTs for the *visible* choices in the *also-critical* condition (42% of the trials) were also significantly slower than in the *also-true* condition and there was a significant interaction between these and the *visible* choices in the *stop-critical* and *stop-true* conditions in the same block.

2.3 Summary

The first point to highlight in summarizing our results is a clear and consistent contrast between *stop* and *also*, such that the presupposition of the former is more likely to be taken into consideration in assessing an *exactly one* claim. However, the results for *stop* were somewhat mixed, in that the visible and covered pictures were chosen about equally often in experiment 1. Experiment 2 suggests some variation between blocks as well. Furthermore, the block design yielded a significant increase in visible picture choices in the second block for *also* as well. In terms of reaction time data, we found relative slow-downs for visible pictures choices for *stop* (in the first experiment and the first block of the second experiment), as well as for *also* (in the second block of the second experiment).

3. Discussion

We now turn to a brief discussion of the results in the broader theoretical context, first with respect to the entailment-contrast proposal by Sudo (2012), and next in light of a possible alternative perspective based on local accommodation.

3.1 Entailed vs. Non-entailed Presuppositions?

Prima facie, the contrast between *stop* and *also* seems to support Sudo's approach: if *stop* entails its presupposition while *also* does not, we expect the presupposition of the former, but not the latter, to play a role in assessing *exactly one* statements. However, there are several aspects of the results that do not straightforwardly fit with this view. First, if the presupposed content of *stop* were an integral part of the entailed content, then we would expect participants to consistently choose the visible picture, as it should be straightforwardly compatible with the corresponding sentences. But in experiment 1, participants chose the covered picture about half of the time. To maintain the entailment-contrast hypothesis despite this data, one could appeal to a task-effect that somehow leads people to respond in a way that is actually inconsistent with the sentence's interpretation. For example, one could imagine participants adapting an economic strategy where they only pay attention to the information displayed on Wednesdays (this suffices in all items with other conditions). While this shortcut results in ignoring crucial information on the *stop*-trials on this view, subjects might simply not be paying enough attention to notice. The reaction time results from experiment 1 could be understood in light of this possibility, as covered box choices for the critical *stop* condition were much faster than responses in most other conditions.

Indeed, experiment 2 in part aimed to further assess this hypothesis: having *stop* trials mixed in with *also* (as well as *her*) trials could have been what led subjects to only pay attention to Wednesdays, as the latter don't seem to involve consultation of earlier days. The block design allowed us to assess *stop* trials on their own. The results were more in line with an analysis of *stop* where its presupposition is also entailed, with visible acceptance rates close to ceiling (though still lower than in the true control condition in the first block). However, two problematic points remain for the entailment contrast approach: first, both in the first block of experiment 2 and in experiment 1, visible picture choices for *stop* exhibited a significant slow-down, for which there is no immediate explanation from this perspective – if the presupposition were genuinely included in the entailed content as well, choice of the visible picture should be straightforward. Secondly, and perhaps even more tellingly, the second block of experiment 2 yielded a surprising increase in visible picture choices for *also*. The theoretically most plausible analysis of such responses is that they are based on an interpretation involving local accommodation (in the form proposed by Heim 1983, or one of the technical variations in the literature), yielding the equivalent of (9b). The observed slow-down in RTs is fully in line with this interpretation, given various recent results suggesting that local-accommodation based responses lead to reaction time delays (Chemla & Bott 2013, Romoli & Schwarz 2015). The question then arises what brings about this increase in local accommodation interpretations for *also* in the second block.

3.2 Variation in Local Accommodation?

Given the challenges for the entailment-contrast proposal just reviewed, it seems worthwhile to consider an alternative perspective on the data. In particular, we'd like to put forth the alternative possibility that triggers in general do not contribute their presupposition directly to the entailed content as well, and that all cases where a presupposition seems to

Entailed vs. Non-Entailed Presuppositions

feature as part of the entailed content result from local accommodation. Furthermore, such an approach has to assume that triggers vary in how available this process is for them. While observations along these lines are commonly made in the literature, it is by no means clear how to explain such variation (see our brief comments below). Let us review how such an approach fares in light of the various aspects of our data.

First, the slow-down in RTs for visible picture choices in *stop-critical* is unsurprising, and in line with the parallel effect for *also* in the second block of experiment 2: local accommodation responses have quite generally been found to lead to such slow-downs. (The absence of a slow-down in the second block of experiment 2 could be attributed to a practice effect.) Conversely, note that the slow-down in block 1 could not be explained as effortful avoidance of an economic task strategy that only considers Wednesday events, as there is no motivation to ignore earlier days when only *stop*-trials have been encountered.

Secondly, the mixed results for *stop*, with an even split between visible and covered picture choices in experiment 1 and a slight increase from first to second block in experiment 2 are entirely unsurprising if an effortful mechanism is required to yield a reading consistent with the visible picture. The higher levels of visible picture acceptances in experiment 2 require further explanation, but likely could be attributed to a reduction in task difficulty based on presenting blocks with homogeneous items in terms of the triggers encountered. Be this as it may, it certainly seems more straightforward to explain variation in accessibility of a reading resulting from a relatively costly process, compared to explaining how a hard-coded entailment reading somehow gets overridden by task-specific effects.

Finally, the increase in *visible* picture choices in *also-critical* trials in the second block of experiment 2 receives an interesting explanation on this account as well: frequent access of local accommodation readings for *stop* in the first block seems to prime access to parallel readings for *also* in the second block. Note that this would provide strong evidence of a general mechanism of local accommodation, given the effect across triggers. Furthermore, the lack of priming in the other direction (where the frequent covered picture choices for *also* in the first block would lead to more covered picture choices for *stop* in the second block) is in line with this as well, as the *absence of a process* on one set of trials would not necessarily be expected to prime absence of the same process elsewhere.

All in all, then, a perspective on the data based entirely on local accommodation seems promising. What remains to be fleshed out is the more general question of just why local accommodation should be easier for some triggers (e.g., *stop*) than others (e.g., *also*). This question has of course featured prominently in the rich theoretical literature on distinctions between triggers. But note that at least the perhaps most prominent existing theoretical accounts in terms of the hard vs. soft distinction are incompatible with the notion that no trigger contributes its presupposition to the entailed content as well. As noted above, they all share the assumption that what is commonly taken to be the presupposition of *soft* triggers is lexically introduced as part of the entailed content, and then receives special status through pragmatic reasoning. So fleshing out the perspective in this section theoretically will require finding alternative accounts of variation in availability of local accommodation across triggers. One common notion that may be of help is that some triggers are explicitly anaphoric in ways that others are not. Exploring the viability of an explanation in these terms requires further theoretical consideration, as well as broader empirical investigation

to assess how well such theoretically motivated delineations of triggers line-up with experimental results from paradigms like the one employed here.

Jérémy Zehr, Florian Schwarz
jeremyz@sas.upenn.edu, florians@sas.upenn.edu

References

- Abusch, Dorit. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27:37–80.
- Chemla, Emmanuel. 2009. Presuppositions of quantified sentences: experimental data. *Natural Language Semantics* 17:299–340.
- Chemla, Emmanuel, & Lewis Bott. 2013. Processing presuppositions: Dynamic semantics vs pragmatic enrichment. *Language and Cognitive Processes* 38:241–260.
- Chierchia, Gennaro. 2015. Factivity meets polarity. Berkeley Linguistics Colloquium.
- Gajewski, Jon. 2016. Another look at npis in definite descriptions: An experimental approach. In *Negation and polarity: Experimental perspectives*, ed. P. Larrivé & C. Lee, Language, Cognition, and Mind 1, 307–327. Springer International Publishing.
- Gajewski, Jon, & I-ta Chris Hsieh. 2014. Comments on negative polarity items in definite description. In *The art and craft of semantics: A Festschrift for Irene Heim*, ed. Luka Crnic & Uli Sauerland, volume 1, 181–198. MITWPL.
- Heim, Irene. 1983. On the projection problem for presuppositions. In *Proceedings of WCCFL 2*, ed. M. Barlow, D. Flickinger, & N. Wiegand, 114–125. Stanford University.
- Huang, Y., E. Spelke, & J. Snedeker. 2013. What exactly do number words mean? *Language Learning and Development* 9:105–129.
- Romoli, Jacopo. 2014. The presuppositions of soft triggers are obligatory scalar implicatures. *Journal of Semantics* (Advance Access).
- Romoli, Jacopo, & Florian Schwarz. 2015. An experimental comparison between presuppositions and indirect scalar implicatures. In *Experimental perspectives on presuppositions*, ed. Florian Schwarz, 215–240. Cham: Springer International Publishing.
- Simons, Mandy. 2001. On the conversational basis of some presuppositions. In *Proceedings of SALT 11*, ed. Rachel Hastings, Brendan Jackson, & Zsófia Zvolenszky, 431–448. Ithaca, NY: CLC Publications.
- Sudo, Yasutada. 2012. On the semantics of phi features on pronouns. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Sudo, Yasutada, & Girosos Spathas. 2015. Gendered nouns and nominal ellipsis in Greek. Ms., UCL.
- Tonhauser, J., D. Beaver, C. Roberts, & M. Simons. 2013. Towards a taxonomy of projective content. *Language* 89:66–109.
- Zeevat, Henk. 1992. Presupposition and accommodation in update semantics. *Journal of Semantics* 9:379–412.