Every restrictors also filter extraposed, that filter canonically[*]

Alexandros Kalomoiros
Heinrich Heine Universität Düsseldorf

Paul Marty
Universidade de Lisboa

Florian Schwarz
University of Pennsylvania

Jacopo Romoli
Heinrich Heine Universität Düsseldorf

Abstract. We experimentally investigate the effects of linear order on presupposition filtering in quantificational constructions. Specifically, we test whether extraposing part of the restrictor of a quantifier like *every* affects how presuppositions in the scope are filtered. Our finding is that extraposition does not affect filtering. This challenges theories that assume a tight correspondence between linear order and filtering, and favours theories that posit a looser connection. We review alternatives that satisfy this desideratum, and outline ways future work can distinguish among them.

Keywords: presupposition filtering, linear order, quantifiers, extraposition

---

* Acknowledgments: To be added...

1 Introduction

Does *linear order* have a role to play in the mechanism that regulates presupposition filtering? This question has been central to approaches that tackle presupposition projection ever since Karttunen and Stalnaker observed that the answer appears to be 'yes' when it comes to conjunction (Karttunen 1973, 1974, Stalnaker 1974):

(1) a. #John stopped smoking and he used to smoke.
b. ✓John used to smoke and he stopped smoking.

This 'yes' was recently put to the test and confirmed in a series of experiments in Mandelkern et al. (2020). Nevertheless, subsequent work complicated the picture: Kalomoiros & Schwarz (2024) replicated the effect of order for conjunction, but showed that disjunction behaves differently: the corresponding disjunctive cases, exemplified below, did not show any order effect.[1]

(2) a. ✓Either Cynthia has stopped raising bees or she has never raised any bees.
b. ✓Either Cynthia has never raised any bees or she stopped raising bees.

Therefore, the effect of order on filtering is not monolithic, but rather modulated by features of the environment surrounding a presupposition trigger, like the connective involved.

The present paper stays with the question of how the effect of order depends on features of the environment, but shifts the focus from connectives to quantifiers. Specifically, building on suggestions in Chierchia 2009, we examine whether restrictor extraposition in English affects filtering. The cases of interest take the following form:

(3) a. Every teacher ***who did cardio yesterday*** exercised again today.
b. Every teacher exercised again today, ***who did cardio yesterday***.

Since the relative clause in the restrictor entails the presupposition of the scope, this should lead to garden-variety filtering in (3a). But in the extraposed version, (3b), the relative clause might come 'too late' to be taken into account, thus leading to the presupposition *Every teacher exercised before today* being projected.[2]

1 Thus experimentally confirming an intuition put forward at least since Karttunen 1974, which however had remained controversial.

2 A similar case involves conditionals with pre- vs. post-posed antecedents, (ia) vs. (ib) (Heim 1990, Schlenker 2009 among others), since antecedents can be analyzed as restrictors (Kratzer 1986). Given the greater controversy surrounding conditional semantics, we leave these for future work.

(i) a. If John ***did cardio yesterday***, he exercised again today.
b. John exercised again today, if ***he did cardio yesterday***.

Whether the extraposed relative indeed comes 'too late' varies by theory: strict linear order approaches expect (3b) to project. But, as we will see, more flexible notions are also possible, which predict filtering for both (3a) and (3b).

There is an intuition that (3b) is slightly degraded, which might seem to show the presupposition projects in (3b). But this is too quick: the degradation could stem from projection or from extraposition itself. Thus, resolving the status of (3b) requires baselines separating the effects of extraposition from those of the presupposition trigger – motivating our controlled experiments.

The rest of the paper is organised as follows. Section 2 examines three views on order and their predictions for (3). Section 3 presents two experiments testing the empirical picture: the result is that extraposition has no effect: both (3a) and (3b) filter the presupposition of the scope equally. Section 4 discusses the implications of this result and concludes.

2 Three Frameworks

Broadly speaking, there are three views on what notion of order matters for presupposition filtering:

- Only linear order matters. (**Strictly Linear view**)
- Only the order of semantic composition matters. (**Meaning-centered view**)
- Linear order matters, but this is modulated by truth conditional features of the environment a trigger finds itself in. (**Intermediate view**)

As we will see in this section, meaning-centered and intermediate approaches can be grouped together when examining our extraposition cases: they both predict filtering for (3), in contrast to the strictly linear view. Nevertheless, it is still useful to differentiate them: as we argue in section 4, they diverge when we look at the conjunction vs. disjunction contrast.

2.1 Strictly Linear

The *strictly linear* view is the natural theoretical stance once the existence of order effects for filtering is taken into account. In its modern form, it boils down to the following two claims (Schlenker 2007, Schlenker 2008, 2009, Chemla & Schlenker 2012, Rothschild 2011, with important antecedents in Karttunen 1974, Stalnaker 1974 and Heim 1983):

- Presuppositions are preferentially filtered by material that linearly precedes them.
- Material that follows a presupposition is harder to use for filtering computations. But access to it is not impossible: it can happen at a processing cost.

These assumptions are formalized in two algorithms: an asymmetric one using only preceding material, and a symmetric one using the whole sentence. Costly access to material following the presupposition is captured by a condition that using the symmetric algorithm comes with processing difficulties.

The predictions for our extraposition paradigm are straightforward: normal asymmetric filtering occurs in (3a), but only the costly symmetric version is available in (3b), which should thus show a presupposition-specific cost over and above any extraposition cost.

2.2 Meaning-Centered

The meaning-centered view takes *order of meaning composition* as the relevant notion of order underlying filtering. The clearest exposition of this is Chierchia 2009 who defines a notion of f-command that holds between two co-arguments of a functor:

(4) Given a functor f that takes two arguments, A and B,

- A ***f-commands*** B iff the first functional complex $f(A)$ that contains A does not contain B

In other words, A f-commands B if and only if A composes with f first creating a complex that leaves B out.

Then, given a functional complex $f(A)(B)$, A is defined as the local context of B just in case A f-commands B. For a presupposition to be filtered then, it must be entailed by its local context, i.e. by material that f-commands it. Chierchia (2009) formalizes this framework within a dynamic system (Heim 1983: a.o.). The exact details of this do not matter here. It's enough to figure out the order in which constituents combine in (3a) vs (3b).

To do this, we need to import an analysis of extraposed constituents. For concreteness, we adopt the account in Fox & Nissenbaum 1999 and analyze the relevant sentences as below (strikeouts represent that something is not pronounced; the results of trace-conversion (Fox 1999) are indicated as bracketed definites):

(5)

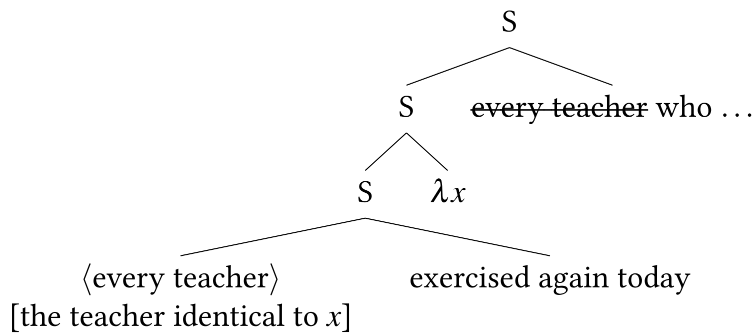

The take-away is that *every* combines with *teacher who did cardio yesterday* first, excluding the scope, *exercised again today*. Note that the lower *every* is pronounced but crucially not interpreted quantificationally (rather trace-conversion applied to it). The higher *every* carries the quantifier denotation, composing with *teacher who did cardio yesterday*, then with the lower λ-abstracted predicate.

Therefore, *teacher who did cardio yesterday* f-commands *exercised again today* serving as its local context; and clearly *teacher who did cardio yesterday* entails the presupposition of *exercised again today*. Thus, filtering is expected. This also holds for the non-extraposed case, since the order of meaning composition remains identical.

2.3 Intermediate

The intermediate approach tries to strike a balance: linear order has a role to play, but the semantic environment also needs to be considered. It comes in two instantiations: the trivalent approach of *Disappointment* (George 2008), and the bivalent approach of *Limited Symmetry* (Kalomoiros 2023b). Both of these theories were designed to account for the conjunction vs disjunction contrast mentioned in section 1. Again, we focus mostly on extraposition, postponing analysis of how these approaches further diverge until section 4.

2.3.1 Disappointment

Disappointment (George 2008) is a trivalent approach to presupposition where connectives and quantifiers are assumed to denote functors that take their arguments in a way that mirrors the linear order in which these arguments appear. For example, *and* denotes a functor $\wedge$ that combines first with the left conjunct, A, and subsequently with the right conjunct B. Thus, the theory assumes a level of representation where order of composition reflects linear order, and is not read off the classical LF typically given to a conjunction, where *and* combines with B first, followed by A.

Being trivalent, the functors can take arguments that are, or return, the # value. #-arguments are dealt with by seeing what happens when they are 'repaired' by classical values. For example: if $A = \#$ and $B = 0$, then $\wedge(A,B)$ is determined by considering repairing $A = \#$ with $A = 1$ and $A = 0$:

- $\wedge(1,0) = 0$
- $\wedge(0,0) = 0$

Here, the result is always 0, so $\wedge(A,B)$ is assigned the value 0. Or at least this is the case on the classic Strong Kleene algorithm for dealing with trivalence. The twist that *Disappointment* adds is that before trying to see what happens under different repairs of #, one checks whether it's possible to get to truth (under the Strong Kleene logic) in the first place; if that's not possible, then # is *disappointing*, and it projects.

Consider conjunction again: if $A = \#$, there is no way to find a B such that $\wedge(A,B) = 1$. Ergo, $\wedge(A,B) = \#$, for any B.[3]

A first disjunct on the other hand is never disappointing: $(\# \vee B)$ always has hope for a truth: just choose $B = 1$. So the normal symmetric Strong Kleene logic applies in this case, in contrast to conjunction.

Extending this to quantifiers requires adapting the linear-order intuition to handle extraposed restrictors: there are some difficulties associated with this, but for the purposes of our paper we can view the extraposed relative as a third argument (C), alongside the restrictor (A) and the scope (B) arguments: *every(A)(B)(C)*.

The details of this move are in Appendix A, where we apply *Disappointment* to this kind of structure. The upshot is that even when the scope B carries a presupposition (and hence is undefined for some entities that are true of the restrictor A) this is not fatal: the extraposed argument C can still restrict A to just those A-entities on which B is fully defined.

The prediction then for our cases of interest is that *Disappointment* expects filtering in both the canonical and extraposed cases, just like the meaning-centered approach.

We now show that the same result obtains for the other instantiation of the intermediate approach.

3 Technically, in $\wedge(A,B)$ where $A = \#$, *Disappointment* checks whether for all A' that are presupposition-equivalent to A, there is B that can make the sentence true. Ignoring this is harmless for propositional cases, but it becomes important when dealing with quantifiers, see Appendix A.

2.3.2 Limited Symmetry

Limited Symmetry (Kalomoiros 2023b,a, Kalomoiros & Schwarz 2024) builds on Schlenker's *Transparency Theory* (Schlenker 2007, Schlenker 2008). Accordingly, it assumes a fully bivalent system where presuppositional sentences are represented in the (formal) language as $p'p$ where p' is the presupposition and p the assertion; this is interpreted conjunctively $p'p = (p' \wedge p)$.

Given a sentence S that embeds a presuppositional expression $p'p$ the goal is to show that for any assertive component p, S is equivalent to S' where S' is the version of S where $p'p$ has been replaced by p.[4] In other words, presuppositional components must be removable from S *salva veritate* regardless of the assertive component p. This is a way of getting at the Stalnakerian intuition that presuppositions are redundant information. Of course, what was just described in the preceding sentences is Schlenker's Transparency constraint; on its own, it will deliver fully symmetric filtering conditions across the board.

The twist *Limited Symmetry* offers is an algorithm about how comprehenders go about establishing the Transparency constraint as they are parsing sentences from left to right.

The algorithm assumes comprehenders parse a sentence while updating the context, trying at each point to compute worlds where the sentence is already true (or false) regardless of continuation. If the part of S they are parsing contains a presuppositional component $p'p$, they also try to establish whether the worlds they are sure S is already true (or false) are worlds where S' (the version of the S where $p'p$ is replaced by p) is also true (or false), for all p. If, following this process of parsing and checking, they reach the end of sentence without any check failing, they will have established that the Transparency constraint holds. Filtering (a-)symmetries are thus a by-product of how this checking algorithm can fail mid-parse, depending on the truth-conditional environment.

To make this more explicit, we apply the algorithm to our cases of interest. This is a novel extension, as Kalomoiros (2023b) doesn't examine quantifiers.

First, we introduce the very simple formal language below:

- **Atomic predicates:** $P, Q, R \ldots$
- **Presuppositional predicates:** $P'P, Q'Q, R'R \ldots$

4 For simplicity, we assume that S contains only one instance of $p'p$. This is not a real limitation; any sentence S which embeds multiple instances of $p'p$ can be re-written by replacing each of the extra instances by sentences of the form $q'p$, $r'p$, ... with the stipulation that $p', q', r' \ldots$ are all equivalent.

- **Formulas:** $Every(A)(B) \mid Every(A)(B)$ who C (where A, B, C are either atomic or presuppositional predicates)

The semantics for this language is the obvious one, with atomic predicates understood as $\langle s, \langle et \rangle\rangle$ functions, presuppositional predicates as $\langle s, \langle et \rangle\rangle$ functions formed by taking the (generalized) conjunction of the P' and P parts; the quantifier formulas are interpreted as below:

(6) a. $[\![\textbf{Every(A)(B)}]\!] = \lambda w.\forall x[A(w)(x) \rightarrow B(w)(x)]$
b. $[\![\textbf{Every(A)(B) who C}]\!] = \lambda w.\forall x[A(w)(x) \wedge C(w)(x) \rightarrow B(w)(x)]$

Let's apply the algorithm to $Every(T)(E'E)$ *who C* (which should be understood as formalizing our extraposed case). The idea is that comprehenders parse this from left to right, symbol by symbol. At each such parsing point, they are trying to figure out worlds in the context c where the sentence is already true/false. The sequence of parsing steps looks like this:

(7) [$Every$, $Every($, $Every(T$, $Every(T)$, $Every(T)($, $Every(T)(E'E$, $Every(T)(E'E)$, $Every(T)(E'E)$ who, $Every(T)(E'E)$ who C]

So for example, at parsing step $Every(T)$, a comprehender knows that all worlds in the context where T denotes the empty set are worlds where the sentence is true no matter how it continues: it doesn't matter what the scope is, or if a further extraposed restriction is coming; in all of these cases the restrictor is empty and the sentence is trivially true.

The constraint on presuppositions will play a non-trivial role as comprehenders encounter a $P'P$-type constituent. In our case, this is at parsing point $Every(T)(E'E)$. At this point, we know the following:

- The sentence is already true in worlds where every T-entity is an $E'E$-entity. Even if a continuation of the form `who C` is encountered, this will just restrict the T entities, but since all of the T entities are $E'E$ entities, this will not change the truth conditions.

- In worlds where $E'E$ is false of at least some of the T-entities, the value of the sentence cannot yet be determined. It might be that the relative clause coming up restricts T in a way that now makes $E'E$ true of all its entities (e.g. in the extreme case, a continuation of the form `who` $\perp$, where $\perp$ is a contradictory predicate, creating the empty set when conjoined with T, so that the whole restrictor ends up empty and the sentence trivially true). But it might also be that just the right *who*-clause comes up that doesn't restrict T to just the $E'E$ entities (e.g. in the extreme case on the other side, a continuation of the form who $\top$, where $\top$ is a tautological predicate), and the sentence ends up false.

The *Limited Symmetry* twist is that we now need to check for the two classes of worlds above where truth value has already been determined, whether they are worlds where the truth value of the sentence without presupposition has already been determined at the corresponding point, for all E. So, in the case of the scope, we look at $Every(T)(E)$ where E' has been removed:

- $Every(T)(E)$ is already true in worlds where every T-entity is an E-entity. The reasoning is very similar as above.
- In worlds where E is false of at least some of the T-entities, the value of the sentence cannot yet be determined, again as before.

Note that all of the worlds where every T-entity is an $E'E$ entity are worlds where every T-entity is an E since $E'E$ is stronger than E. This bit of reasoning doesn't depend on E, so it holds for all E. So, for any E the worlds where the sentence is already true at parsing point $Every(T)(E'E)$ are worlds where the presuppositionless version is also already true (at the corresponding parsing point). The constraint then is satisfied and the parse moves on.

The next bit of action happens when the entire extraposed relative has been parsed. At that point, there is nothing more that can follow, given the syntax we specified, and the comprehender needs to check whether:

- For all E, all the worlds where $Every(T)(E'E)$ who C is true are worlds where $Every(T)(E)$ who C is true.
- For all E, all the worlds where $Every(T)(E'E)$ who C is false are worlds where $Every(T)(E)$ who C is false.

Expressed more symbolically, these requirements can be written as:

- For all E, $\{w \in c|$ every $T \wedge C$-entity is an $E'E$-entity$\} \subseteq$ $\{w \in c|$ every $T \wedge C$-entity is an E-entity$\}$
- For all E, $\{w \in c|$ some $T \wedge C$-entity is not an $E'E$-entity$\} \subseteq$ $\{w \in c|$ some $T \wedge C$-entity is not an E-entity$\}$

Suppose that in a world, every $T \wedge C$-entity is an $E'E$-entity. Then *a fortiori* every $T \wedge C$-entity is an E-entity. This doesn't depend on E so it holds for all E.

Suppose that in a world, some $T \wedge C$-entity is not an $E'E$-entity. So, it's either not an E'-entity or not an E-entity. But recall that in the case of interest, C (*who did cardio yesterday*) entails the presupposition of the scope (*exercised again today*). In our formalization this means that $C \models E'$, hence $T \wedge C \models E'$. So, in our world, the $T \wedge C$ entities that are not $E'E$ must be not E entities. Therefore, our world is also a world where $Every(T)(E)$ who C is false. Once again this did not depend on E, so it goes through for all E.[5]

The upshot of these little proofs is that the *Limited Symmetry* version of the *Transparency* constraint does not impose any requirement on the context in this case: in parsing $Every(T)(E'E)$ who C, the relevant requirements are satisfied at each step and the end is reached with no issue; hence, the presupposition of the scope gets filtered.

The same holds for the non-extraposed version $Every(T \wedge C)(E'E)$, as the reader is invited to verify.

Summing up, *Limited Symmetry*, just like *Disappointment*, expects no presupposition-specific difference between the extraposed and the non-extraposed version of sentences like (3).

3 The Experiments

To test whether extraposing part of the restrictor leads to projection of presuppositions in the scope we ran two studies. Their preregistrations, together with lists of the experimental items and R scripts for the analyses reported below, are accessible on the paper's Zenodo page.

3.1 Design

We want to test whether (8b) shows projection, in contrast to (8a):

(8) a. Every teacher who did cardio yesterday exercised again today. (PsSecond)
    b. Every teacher exercised again today, who did cardio yesterday. (PsFirst)

The labels PsSecond/First refer to whether the presupposition comes before or after the part of the restrictor that can filter it: before in PsFirst, after in PsSecond.

As mentioned in section 1, we need to separate the effect of extraposition from any effects of presupposition. This is done via the following baselines:

5 In general, *Limited Symmetry* requires that a sentence like *Every(R)(Q'Q) who L* presupposes that *for every world w in the context c, for every entity x in w:* $(R(w)(x) \wedge L(w)(x)) \rightarrow Q'(w)(x)$.

(9) ***Support:*** At our school, there was a new initiative to include exercise in teachers' schedules. The schedule comes in two versions: the first includes exercise on only one day, the second is spread out across multiple days. The initiative just started yesterday. **I know that all the teachers exercised yesterday**, and some of them did cardio. Today, I learned that:

a. Every teacher who did cardio yesterday exercised **again** today. (**S-PsSecond**)

b. Every teacher exercised **again** today, who did cardio yesterday. (**S-PsFirst**)

The idea is that in the S(upport) contexts, the context makes the presupposition projected in PsFirst (under the strictly linear approach) true: it affirms that *all the teachers exercised yesterday and some of them did cardio* (the *cardio* part ensures that the restrictor isn't empty).

Thus filtering plays no role in the Support context: the presupposition is already supported for both PsFirst and PsSecond. So any difference between PsFirst vs PsSecond can wholly be attributed to the extraposition that has taken place in PsFirst but not in PsSecond; in other words, these conditions control for the cost of extraposition.

To test whether the sentences carry an extra cost when comprehenders have to check whether the presupposition gets filtered, we utilize so-called Explicit Ignorance (EI) contexts (Simons (2001), Abusch (2010); see also Mandelkern et al. (2020), Kalomoiros & Schwarz (2024)):

(10) ***Explicit Ignorance:*** At our school, there was a new initiative to include exercise in teachers' schedules. The schedule comes in two versions: the first includes exercise on only one day, the second is spread out across multiple days. The initiative just started yesterday. I know that some of the teachers did cardio yesterday, although **I have no idea if every teacher exercised in some way**. Today, I learned that:

a. Every teacher who did cardio yesterday exercised **again** today. (**EI-PsSecond**)

b. Every teacher exercised **again** today, who did cardio yesterday. (**EI-PsFirst**)

The idea is simple: (10a) projects no presupposition and is fine in any context. Conversely, (10b) projects (at least under the strict linear view) that *every teacher exercised before today*. This is explicitly in doubt in the EI context; thus, in addition to any costs incurred by extraposition, (10b) should also reflect costs associated with projecting a presupposition in a context that doesn't support it.

The strictly linear view expects the difference in acceptability between EI-PsFirst vs EI-PsSecond to reflect both projection and extraposition costs, and thus be larger than the difference between S-PsFirst vs S-PsSecond, which reflects only extraposition costs. In other words an interaction between Context (EI vs S) and Order (First vs Second) is predicted.

The meaning-centered and intermediate approaches predict no such interaction: since costless garden-variety filtering is predicted for both EI-PsFirst and EI-PsSecond, the only difference between them should be the cost of extraposition, already measured in the S contexts.

Returning to the EI conditions, note that if a presupposition *every teacher exercised before today* is projected, it's not something that can be globally accommodated: that would produce a contradiction in the context. Thus the only options for such a presupposition are for it to be filtered somehow (e.g. via some form of symmetric filtering) or to be locally accommodated (cf. Heim (1983)).

To test whether comprehenders would opt to locally accommodate such a presupposition, the design included an extra baseline, dubbed *NoFilt*:

(11) a. ***EI:****...some of the teachers exercised; but* ***I don't know if every teacher exercised in some way****...*

b. ***S:****...***and I know that every teacher exercised***...*
Today, I learned that:

c. Every teacher exercised **again** today. (**NoFilt**)

NoFilt are sentences where any material that can filter the presupposition of the scope has been removed: under all theories they are expected to project *every teacher exercised before today*. This is satisfied in the S context, but in explicit doubt in the EI context. In the latter context the only option is local accommodation (since it cannot be globally accommodated as this would lead to a contradiction; and filtering isn't an option either given that there is no linguistic material capable of doing this).

Thus, the difference between EI-NoFilt vs S-NoFilt will give us an estimate of the cost of local accommodation. This then can be compared to any difference between EI-PsFirst vs EI-PsSecond: if that difference is less than the corresponding difference in the NoFilt conditions, this is evidence that some form of symmetric filtering is taking place in the EI-Ps conditions.[6]

6 We assume throughout that symmetric filtering, however costly, remains cheaper than local accommodation; otherwise it would play no role in any of the theories under discussion.

All theories then expect the following interaction: S-NoFilt should match EI-PsSecond since in both the presupposition receives prior support (either contextually or through costless filtering); conversely, EI-NoFilt should be less acceptable than EI-PsFirst: the former requires local accommodation, whereas the latter involves symmetric filtering that on any theory is less costly than local accommodation. We will call this interaction the Filt $\times$ PriorS interaction (see section 3.3.)

In total we created 24 items like the one illustrated above, for each of these 6 conditions. We also constructed 24 fillers.[7]

The items in both studies were similar, with one difference: Study 1 used *again, too, re-* (*re-* as in re-examine etc.) as presupposition triggers, whereas Study 2 used *also* in the place of *too.*

3.2 Procedure

We made Context a between-subjects factor; as such, Group 1 of participants saw the EI-Ps and EI-NoFilt conditions, while Group 2 saw the S-Ps and S-NoFilt conditions.

For Study 1, we recruited 200 participants (100 per between-subjects group) whereas for Study 2 300 (150 per group). Across the three conditions of each between-subjects group, items were counterbalanced in a Latin Square design.

Participants were recruited on Prolific and screened for native English speakers; we also preregistered an exclusion criterion based on our filler items.[8] Recruitment continued until we reached the preregistered number of participants in each study post exclusions.

The task was an acceptability judgment: participants saw a context, then a sentence, and rated its naturalness on a 7-point scale. The Support and Explicit Ignorance parts were bolded (see Fig. 1). All target items were presented in random order.

Before seeing any of our target sentences, participants went through three training trials. These consisted of three contexts and sentences: one was unproblematic, one contradictory and one contained an unsupported implicature. In each case, participants received feedback on their answers.

7 Fillers were of two types: Good/Bad conditionals (differing in whether the antecedent was contextually possible) and True/False sentences. See the Zenodo page for examples.

8 Study 1 excluded participants scoring under 2 points higher on Good than Bad conditionals, or under 4 points higher on True than False fillers; Study 2 retained only the Good/Bad exclusion, as it alone was sufficient in Study 1. Full criteria are in the preregistrations on Zenodo. See also fn. 7.

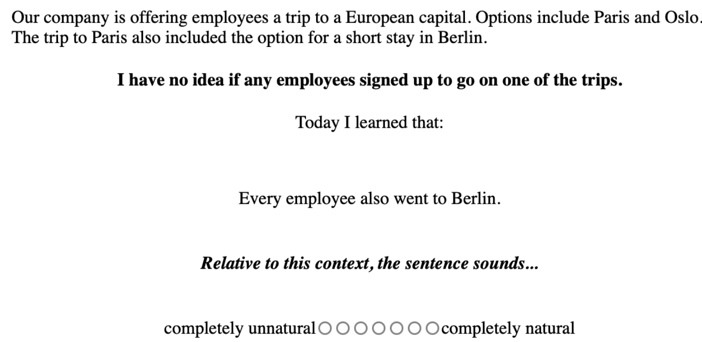

Figure 1 Example Trial

3.3 Results

The results for both studies are depicted in Figure 2. Overall there is no visible effect of a Context $\times$ Order interaction. At the same time, the difference in the NoFilt conditions is larger than the corresponding difference in the EI-Ps conditions, pointing to a Filt $\times$ PriorS interaction.

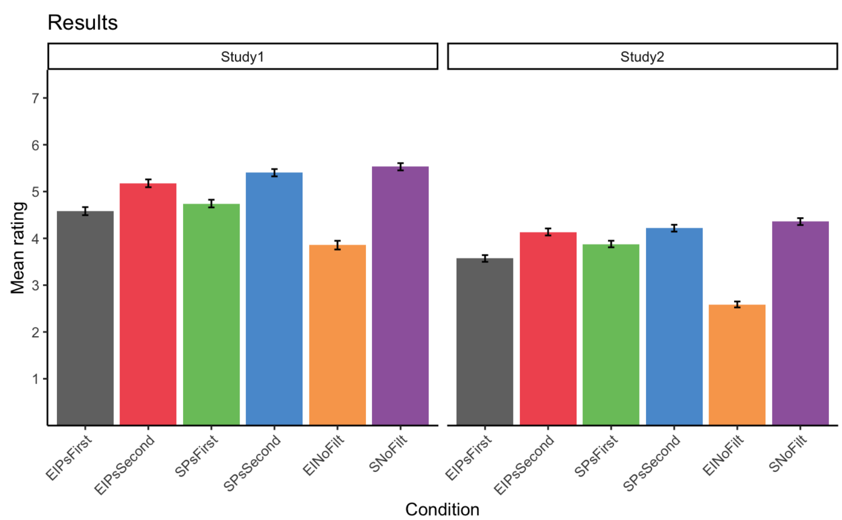

Figure 2 Mean Rating across conditions for Studies 1, 2. Error bars represent standard error.

For Study 1 we preregistered a frequentist analysis of the interactions of interest. The relevant mixed effects models showed the absence of a significant Context $\times$ Order interaction ($\beta = 0.02$, $SE = 0.06$, $z = 0.37$, $p = .7$) and the presence of a Filt

$\times$ PriorS interaction ($\beta = 0.29$, $SE = 0.095$, $z = 3.04$, $p = .002$). Once it became evident that the data leaned towards the absence of Context $\times$ Order interaction, we decided to switch to a Bayesian analysis in the preregistration for Study 2, and also proceeded to reanalyze the Study 1 results in the same framework. For reasons of space we focus on the Bayesian analyses of both studies here, but the corresponding frequentist analyses are available on Zenodo.

We evaluated both interactions by fitting the Bayesian maximal mixed effects model allowed by our design, with a cumulative *logit* likelihood. All factors were sum-coded.

Starting with the Context $\times$ Order interaction, we used *brms* to fit a Bayesian model on the Study 1 data, predicting Response from Context, Order, and their interaction. The model included random intercepts for Participant and Item, as well as by-participant random slopes for Order, and by-item random slopes for Context, Order, and their interaction.[9] We fit the model with default brms priors, then ran a sensitivity analysis under two alternative prior specifications:[10]

- an *agnostic* prior, where all fixed effects were assigned a *normal(0, 1.5)* prior. This distribution allows for both positive and negative effects while remaining centered at 0, expressing neutrality regarding the presence of an effect.

- an *adversarial* prior, which retained the *normal(0, 1.5)* prior for all fixed effects except the Context $\times$ Order interaction, which was assigned a *normal(-0.42, 0.33)* prior. This specification represents a strong skeptical baseline by encoding a prior belief that favors a substantial interaction effect.

The adversarial prior was derived by re-analyzing Experiment 3 from Mandelkern et al. (2020). While that experiment featured a design parallel to ours, it investigated presupposition filtering with conjunctions rather than quantifiers. Because their results indicated a robust Context $\times$ Order interaction, we leveraged their posterior estimates to construct an informative prior representing the expected shape and direction of the interaction, should it exist in our data.

We carried out the same analysis for the Study 2 data, adding another prior in the sensitivity analysis that was derived from the Study 1 posterior (under the default

9 This is the maximal model. Recall that Context was between-subjects, so a by-participant random slope for it cannot sensibly be included.

10 The Study 2 preregistration specified posterior estimates under default and informed priors, a bridge-sampling Bayes Factor just on the Study 2 data, and a cross-dataset trajectory analysis (see Appendix B). We additionally report a sensitivity analysis using agnostic and adversarial priors, not preregistered, which we extended to Study 1 and the combined dataset for consistency.

prior).[11] As a further confirmatory step, we extended the analysis to the combined Study 1+Study 2 dataset, under the agnostic and adversarial priors.[12]

The resulting posterior estimates for the Context × Order interaction are summarized in Table 1. Across the different priors, the posterior is close to 0, indicating that the null hypothesis cannot be excluded (see also the graph in Appendix B that tracks the development of this posterior as the dataset grows). A Bayes Factor analysis then compared each model to its counterpart excluding the interaction: BF_{01} always favoured the null model (Table 1), providing clear evidence against a Context × Order interaction (throughout, Bayes factors are reported only for the sensitivity-analysis priors, not the default prior).[13,14]

Dataset	Prior	Estimate	*SD*	95% CrI		$\hat{R}$	BF_{01}
				Lower	Upper		
Study 1	Default	0.03	0.07	−0.11	0.17	1.00	–
	Agnostic	0.03	0.07	−0.12	0.17	1.00	19.18
	Adversarial	0.00	0.07	−0.13	0.14	1.00	9.64
Study 2	Default	−0.08	0.05	−0.17	0.01	1.00	–
	Agnostic	−0.08	0.05	−0.17	0.01	1.00	7.45
	Adversarial	−0.09	0.05	−0.18	0.00	1.00	2.78
	Informed	−0.08	0.05	−0.17	0.01	1.00	4.34
Study 1+Study 2	Agnostic	−0.04	0.04	−0.12	0.03	1.00	16.51
	Adversarial	−0.05	0.05	−0.12	0.03	1.00	4.85

Table 1 Posterior estimates for the Context × Order interaction across datasets, with Bayes Factors (BF_{01}) comparing models without vs with the interaction term. $\hat{R}$ values of 1.00 indicate model convergence.

Turning to the second interaction (defined in section 3.1), we operationalize it via two factors: **i)** Filt(ering) categorized the relevant conditions as either involving material that can potentially filter the presupposition of the scope (so EI-PsFirst and EI-PsSecond were tagged Filt; whereas the NoFilt conditions were tagged NoFilt) **ii)**

11 We used the means from the Study 1 posterior, but kept the *SD*s to 0.5 to remain regularizing.

12 Including an informed prior derived from Study 1 or 2 would essentially constitute a form of double-counting: the data being tested would already inform the prior.

13 The lower BF_{01} under the adversarial prior reflects the prior's design—placing greater mass on the presence of an interaction—rather than a weakening of evidence; even under this skeptical prior, the null model is preferred.

14 We also conducted a frequentist analysis on Study 2. For the Order × Context interaction, the result was null: $\beta = -0.07$, $SE = 0.04$, $z = -1.76$, $p = .078$.

A PriorS factor that tagged conditions on the basis of whether the presupposition received prior support of some kind (EI-PsSecond and S-NoFilt were tagged as PriorS, whereas EI-PsFirst and EI-NoFilt as NoPriorS). The relevant interaction boils down to whether the effect of PriorS is larger in the NoFilt conditions compared to the Filt cases.

We again fit models across default, agnostic, adversarial, and informed priors (the latter for Study 2 only); the resulting posteriors excluded 0 (except on the adversarial prior for Study 1, see also fn. 17), pointing to a genuine effect (see also Appendix B).[15,16] This was further corroborated by the BF analyses which always favoured the interaction model (see Table 2).[17,18]

Dataset	Prior	Estimate	*SD*	95% CrI		$\hat{R}$	BF_{10}
				Lower	Upper		
Study 1	Default	0.29	0.1	0.1	0.49	1.00	–
	Agnostic	0.29	0.1	0.1	0.49	1.00	4.13
	Adversarial	0.06	0.05	−0.04	0.15	1.00	1.6
Study 2	Default	0.35	0.07	0.21	0.49	1.00	–
	Agnostic	0.35	0.07	0.21	0.48	1.00	> 100
	Adversarial	0.11	0.04	0.02	0.19	1.00	29.51
	Informed	0.29	0.07	0.15	0.43	1.00	> 100
Study 1+Study 2	Agnostic	0.32	0.06	0.19	0.44	1.00	> 100
	Adversarial	0.11	0.04	0.03	0.2	1.00	43.3

Table 2 Posterior estimates for the Filt × PriorS interaction across datasets, with Bayes Factors (BF_{10}) comparing models with vs without the interaction term.

15 The adversarial prior in this case was set to *normal(0, 0.05)*, encoding skepticism for the presence of an interaction.

16 Not all participants see both levels of Filt and PriorS (crucially, Group 2 participants contribute only S-NoFilt to this analysis, which is tagged NoFilt and PriorS, so they see no Filt-tagged and no NoPriorS-tagged conditions). Hence by-participant random slopes for these factors cannot be estimated across the board and are omitted.

17 The comparatively modest BF_{10} under the adversarial prior in Study 1 (Table 2) reflects the prior's skepticism toward the interaction rather than weak evidence in the data itself. The same holds for the CrI (cf. fn. 13).

18 These results were also supported by a frequentist analysis for Study 2: $\beta = 0.35$, $SE = 0.07$, $z = 5.16$, $p < .001$.

4 Discussion

4.1 The main result

The absence of a presupposition-specific difference between the extraposed and canonical orders of our sentences (together with our Bayesian analysis) provides a substantive argument against strictly linear approaches: even though the filtering material comes ***after*** the trigger in our extraposed cases, this incurs no extra cost beyond that of extraposition itself. Moreover, it is clear that filtering is taking place in these constructions: if it were not, local accommodation would be the only recourse, and we would not observe a Filt $\times$ PriorS interaction.

So we are left with two types of approaches: the meaning-centered approach, which gives linear order no role, and the intermediate approach, where it plays a restricted, modulated role. Our results are compatible with both. The question moving forward is whether we can differentiate them further — first between meaning-centered and intermediate approaches, and second between *Disappointment* and *Limited Symmetry*.

4.2 Moving forward

In what follows we argue that:

- It is hard to arrive at an explanatory version of a meaning-centered approach à la Chierchia (2009) when it comes to projection differences between conjunction vs disjunction.
- The predictions of *Disappointment* vs *Limited Symmetry* come apart when we look at other quantifiers: here we focus on the case of *No* vs *Every*.

4.2.1 Meaning-centered versus the intermediate approaches

While the meaning-centered view is compatible with our results, it faces trouble with the contrast between conjunction vs disjunction (recall from Section 1 that conjunctions filter asymmetrically, but disjunctions symmetrically). Consider the typical asymmetric syntactic representations for conjunction/disjunction in (12) (cf. Munn 1993):

(12)

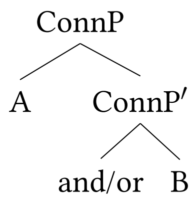

(13)

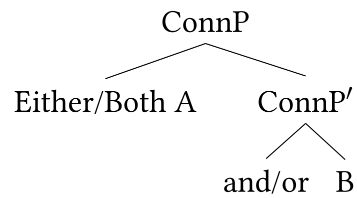

Going by the f-command account of local contexts, the right conjunct/disjunct (B) combines with the functor first, so it f-commands – and forms the local context of – the left conjunct/disjunct (A). The prediction then is that material in B should filter presuppositions in A. Actually, it's even stronger than that: only material in B should contribute to filtering. Since A doesn't f-command B it can never filter presuppositions in it.

This fails for both conjunction and disjunction: presuppositions from the first conjunct project and do not get filtered; on the other hand, the first conjunct/disjunct is always involved in filtering. Chierchia (2009) is aware of this, and suggests a way out by adopting the syntactic representations in (13).

The idea is that it's the (usually silent) both/either elements in these trees that act as the functor carrying the conjunctive and disjunctive meanings respectively: *and/or* are inert and contribute no semantics. Under this view, A always f-commands B and acts as its local context. This move clearly gets the case of conjunction. But it still faces issues with disjunction, since now the expectation is that presuppositions project from the first disjunct and cannot be filtered by the second disjunct: empirically this is problematic.

One fix would be to stipulate that both (12) and (13) are available representations for disjunction, and that filtering happens as long as one disjunct f-commands the other in *at least one* representation. But then the question is why doesn't conjunction show the same representational optionality?

Thus, the meaning-centered approach (at least in Chierchia (2009)) is caught in a bind when it comes to differentiating between conjunction vs disjunction: there doesn't seem to us to be a non-stipulative way of drawing the line between the two.[19]

19 Note that this also holds for strictly linear views (an extra challenge for them, quite apart from the extraposition cases that have been the focus here). See Kalomoiros & Schwarz (2024) for more discussion.

4.2.2 Disappointment versus Limited Symmetry

The two intermediate approaches (*Disappointment* vs. *Limited Symmetry*) are compatible with our experiments, and have the extra advantage of unproblematically handling the conjunction vs disjunction difference. What kind of evidence would differentiate between them?

The answer is that their predictions part when we move to quantifiers beyond *every*. Specifically, *Limited Symmetry* expects *No* to not allow filtering from extraposed material (in contrast to *every*). Conversely, *Disappointment* expects no difference between the two quantifiers.

The more careful derivation for the *Disappointment* prediction can be found in Appendix A. Here we briefly illustrate why *Limited Symmetry* makes a different prediction for *No.*

Consider a version of the cases of interest to us, but with *No* substituted for *Every*:

(14) a. No teacher who did cardio yesterday exercised again today.
b. No teacher exercised again today, who did cardio yesterday.
$\rightsquigarrow No(T)(E'E)$ *who* $\boldsymbol{C}$

The case of (14a) is not very interesting: garden variety filtering is predicted (as the reader can verify). For (14b), the action is at the *No teacher exercised again today* point. It suffices to consider the worlds in the context c where the sentence is already true at this point:

- The sentence is already true in worlds where every T-entity is $\neg E'E$. Even if some extraposed restriction restricts T, since all T falsify $E'E$, the sentence will continue to be true.
- For the sentence without the presupposition, at the corresponding parsing point, we have that $No(T)(E)$ is already true in worlds where every T-entity is $\neg E$.

Limited Symmetry then imposes the following requirement:

(15) For all E: $\{w \in c|$ every T-entity is either not in E or not in E'$\} \subseteq \{w \in c|$ every T-entity is not in E$\}$

This only holds if all worlds in the global context c support that $\forall x : T(x) \rightarrow E'(x)$. For suppose this doesn't hold. Then there is a world $w' \in c$ such that some T-entity is not an E' entity. We can now find an E that falsifies the constraint in (15). Choose $E = \neg E'$. Then the constraint becomes:

(16) $\{w \in c|$ every T-entity is either not in $\neg E'$ or not in $E'\} \subseteq \{w \in c|$ every T-entity is not in $\neg E'\}$

This simplifies to:

(17) $c \subseteq \{w \in c|$ every T-entity is in E'$\}$

But (17) doesn't hold, since we have a world, namely w' where some T-entity is not an E'-entity.

As such the following prediction is made under *Limited Symmetry*:

(18) a. *Every teacher exercised again today, who did cardio yesterday* $\rightsquigarrow$ *presuppositions get filtered*
b. *No teacher exercised again today, who did cardio yesterday* $\rightsquigarrow$ presupposes that *Every teacher exercised before today*

On the other hand, *Disappointment* predicts filtering in both cases.

We would like to end then by suggesting that testing this divergence in predictions is a clear next step for clarifying the interaction between linear order and 'local context'-related phenomena like presupposition filtering.

A *Disappointment* and extraposition

In this appendix, we apply *Disappointment* to extraposed quantificational structures. One difficulty is adapting the intuition that a quantifier's functor combines with its arguments in linear order: the extraposed relative clause is clearly part of the restrictor argument, but appears linearly after the scope.

One option that keeps to the spirit of the intuition is to define a version of *every* that takes three predicate arguments, A, B, C, with the understanding that the C further restricts A ($\wedge$ below represents generalized conjunction):

(19) a. $\text{every}_{norm} = \lambda A_{et}.\lambda B_{et}.\, \forall x[x \in A \rightarrow x \in B]$
b. $\text{every}_{extr} = \lambda A_{et}.\lambda B_{et}.\lambda C_{et}.\, \forall x[x \in A \wedge C \rightarrow x \in B]$

Let's symbolize the cases of interest to us as below, where A can be thought of as *teacher*, C as *who did cardio yesterday* and B as *exercised again today*:

(20) a. $\text{every}_{norm}(A \wedge C)(B)$
b. $\text{every}_{extr}(A)(B)(C)$

We can now reason about how the *every* functor combines with its presuppositional B argument, i.e., B denotes a function from entities to truth values that returns # for at least some entities. In the case where the restrictor only includes entities for which B returns a classical truth value, there is no issue: the sentence is true or false, but never #. Such a case is represented by every$_{norm}(A \wedge C)(B)$, where all the entities in the restrictor satisfy the presupposition of the scope.

For the extraposed case, assume we get the presuppositional scope B, and that for some entities in A, B returns #. To apply *Disappointment* we ask whether every presupposition-equivalent counterpart of B kills off hope of truth, where presupposition equivalence is understood as follows:

(21) B' is presupposition-equivalent to B iff for all $x: \; B'(x) = \# \leftrightarrow B(x) = \#$

That is, a presupposition-equivalent counterpart of B must be undefined on exactly the same entities as B, but is otherwise unconstrained: on any entity where B is classically defined, a counterpart is free to return 1 or 0, independently of what B itself returns there.

This has an important consequence: since a counterpart need not agree with B's actual classical value anywhere, whether B returns 1 or 0 for the entities in A on which it is defined plays no role in whether hope of truth survives. All that matters is whether A contains *some* entity on which B is defined at all. If it does, there is a presupposition-equivalent counterpart of B that assigns 1 to that entity, together with a C that restricts A down to just that entity; under that counterpart, every$_{extr}(A)(B')(C) = 1$, so hope of truth survives no matter what B's actual classical values are. B is disappointing only in the limiting case where B is undefined on *every* entity in A: then every presupposition-equivalent counterpart is undefined throughout A as well, and no choice of C can produce a classical value.

Thus, as long as A contains some entity on which the scope is defined – a property that holds in our contexts, see (10) – *Disappointment* expects no presupposition-specific difference between (3a) vs (3b).

The same reasoning applies to a case like $no_{extr}(A)(B)(C)$: as long as A contains some entity on which B is defined, there is hope of truth – some presupposition-equivalent counterpart assigns 0 to that entity, and C can restrict A down to it – regardless of what B actually returns there or elsewhere in A. Hence no difference is expected between $no_{extr}(A)(B)(C)$ vs $no_{norm}(A \wedge C)(B)$.

B Graphs

As mentioned in section 3.3, we preregistered that we would graph the trajectory of the posteriors for the Context $\times$ Order and Filtering $\times$ PriorS interactions in our studies by looking at how they start out when we restrict attention to just the first 100 participants, then the first 200 and so on (until the full dataset of 500 people is reached for each interaction). The values come from models fit with default brms priors and are graphed below.

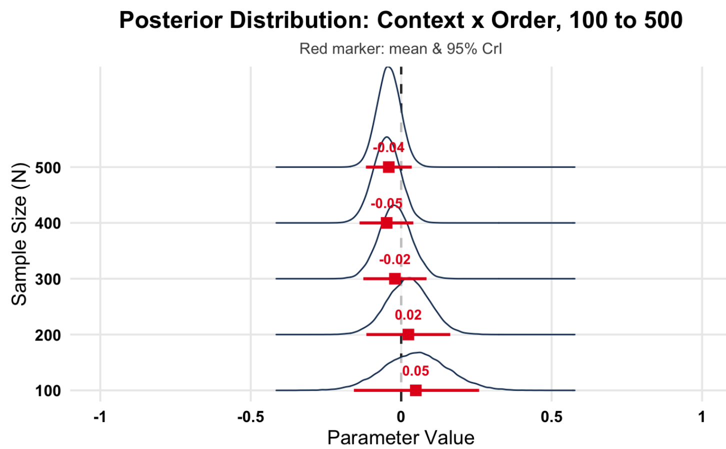

Figure 3 Posterior trajectory for Context × Order as participants are added in increments of 100 (from $n = 100$ to $n = 500$), with 95% credible intervals.

Consistent with the analysis in section 3.3, the posterior for Context × Order hovers around 0, with the 95% CrI never excluding 0, even as more people are added. In contrast, the posterior for Filt × PriorS starts with 0 excluded from the 95% CrI even at the first 100 people, and continues to peak away from 0.

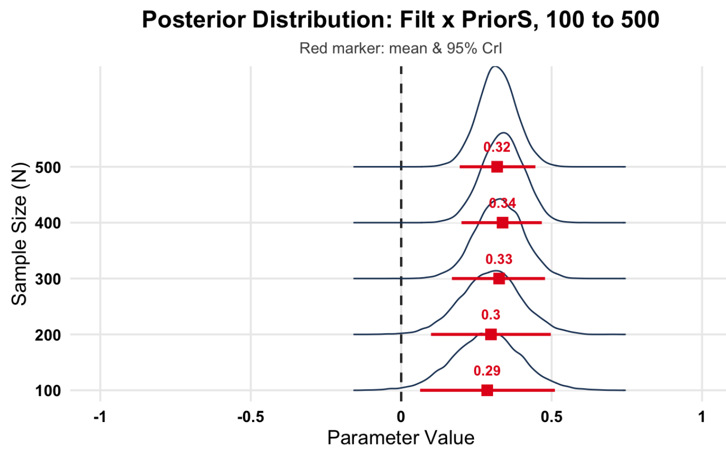

Figure 4 Posterior trajectory for Filt × PriorS as participants are added in increments of 100 (from $n = 100$ to $n = 500$), with 95% credible intervals.

References

Abusch, Dorit. 2010. Presupposition triggering from alternatives. *Journal of Semantics* 27(1). 37–80.

Chemla, Emmanuel & Philippe Schlenker. 2012. Incremental vs. symmetric accounts of presupposition projection: An experimental approach. *Natural Language Semantics* 20(2). 177–226.

Chierchia, Gennaro. 2009. On the explanatory power of dynamic semantics. Handout from talk at Sinn und Bedeutung 14.

Fox, Danny. 1999. Reconstruction, binding theory, and the interpretation of chains. *Linguistic Inquiry* 30(2). 157–196.

Fox, Danny & Jon Nissenbaum. 1999. Extraposition and scope: A case for overt QR. In Sonya Bird, Andrew Carnie, Jason D. Haugen & Peter Norquest (eds.), *Proceedings of the 18th west coast conference on formal linguistics*, 132–144. Somerville, MA: Cascadilla Press.

George, Benjamin R. 2008. A new predictive theory of presupposition projection. In *Semantics and Linguistic Theory*, vol. 18, 358–375.

Heim, Irene. 1983. On the projection problem for presuppositions. In M. Barlow, D. Flickinger & N. Wiegand (eds.), *Proceedings of WCCFL 2*, 114–125.

Heim, Irene. 1990. Presupposition projection. In R. van der Sandt (ed.), *Reader for the Nijmegen workshop on presupposition, lexical meaning, and discourse processes*, University of Nijmegen.

Kalomoiros, Alexandros. 2023a. Deriving presupposition projection in coordinations of polar questions: A reply to Enguehard 2021. *Natural Language*

Semantics 31(4). 253–290.
Kalomoiros, Alexandros. 2023b. *Presupposition and its (a-)symmetries*: University of Pennsylvania dissertation.
Kalomoiros, Alexandros & Florian Schwarz. 2024. Presupposition projection from 'and' vs. 'or': Experimental data and theoretical implications. *Journal of Semantics* 41(3-4). 331–372. http://dx.doi.org/10.1093/jos/ffae011.
Karttunen, Lauri. 1973. Presuppositions of compound sentences. *Linguistic Inquiry* 4(2). 169–193.
Karttunen, Lauri. 1974. Presupposition and linguistic context. *Theoretical Linguistics* 1(1-3). 181–194.
Kratzer, Angelika. 1986. Conditionals. *Chicago Linguistics Society* 22(2). 1–15.
Mandelkern, Matthew, Jérémy Zehr, Jacopo Romoli & Florian Schwarz. 2020. We've discovered that projection across conjunction is asymmetric (and it is!). *Linguistics and Philosophy* 43(5). 473–514.
Munn, Alan Boag. 1993. *Topics in the syntax and semantics of coordinate structures*: University of Maryland, College Park dissertation.
Rothschild, Daniel. 2011. Explaining presupposition projection with dynamic semantics. *Semantics and Pragmatics* 4. 1–43.
Schlenker, Philippe. 2007. Anti-dynamics: Presupposition projection without dynamic semantics. *Journal of logic, language and information* 16. 325–356.
Schlenker, Philippe. 2008. Be articulate: A pragmatic theory of presupposition projection. *Theoretical Linguistics* 157–212.
Schlenker, Philippe. 2009. Local contexts. *Semantics and Pragmatics* 2. 1–78.
Simons, Mandy. 2001. On the conversational basis of some presuppositions. In Rachel Hastings, Brendan Jackson & Zsofia Zvolenszky (eds.), *Semantics and linguistic theory 11*, 431–448. Ithaca, NY: CLC Publications.
Stalnaker, Robert. 1974. Pragmatic presuppositions. In Milton K. Munitz & Peter K. Unger (eds.), *Semantics and Philosophy*, 197–213. New York University Press.