

Clause-internal coherence as presupposition resolution*

Kelsey Sasaki and Daniel Altshuler

University of Oxford, Oxford, United Kingdom
{kelsey.sasaki, daniel.altshuler}@ling-phil.ox.ac.uk

1 Introduction

(1) exemplifies what [12] coins *Clause-Internal Coherence* (CIC): inferences that arise from a single clause, despite being characteristic of (multi-clause) discourse.¹ In particular, a causal inference is salient: by sticking a knife into her husband, the described subject became a widow.

- (1) A widow stuck a knife into her husband. ([3])

Recently, [9] argued that familiar pragmatic tools cannot account for CIC and proposed a previously unidentified type of enrichment, *eliciture*, that is characterized by its non-local character, familiar from research on coherence in intersentential discourse (see [16], [14] for recent overviews). One major outstanding question is how to compositionally derive clause *internal* coherence from clause *external* coherence relations.

This paper takes strides in answering this question. We explore causal inferences with deverbal adjectives, comparing predicative uses in discourse like (2a)-(2b) with attributive, clause-internal uses like (2c)-(2d). §2 offers evidence from offline experiments that (i) attributive deverbals trigger causal inferences, but do so more weakly than their predicative counterparts; (ii) attributive non-deverbals behave similarly to deverbals, but are even weaker causal triggers.

- (2) a. *Discourse effect-cause*: A child was **drenched**. She got hit by a big water balloon.
b. *Discourse cause-effect*: A big water balloon hit a child. She was **drenched**.
c. *Clause effect-cause*: A **drenched** child got hit by a big water balloon.
d. *Clause cause-effect*: A big water balloon hit a **drenched** child.

In §3, we explore the possibility of extending an analysis of presupposition in SDRT ([4]) to explain the experimental findings. This is motivated by the observation that CIC can arise as a presupposition, with deverbal adjectives as presupposition triggers:

- (3) a. It's not the case that a drenched child got hit by a big water balloon. She was pushed into the pool.
b. A few children at Camp Hope showed up to dinner drenched. If a drenched child got hit by a big water balloon, then someone smuggled such balloons into the camp.

The analysis we propose views presupposition as a species of anaphora resolution ([22] and [18]). We show how a causal inference follows from the resolution of a coherence relation that binds the presupposed information and an attachment point that allows for projection. The salience of a causal inference, we argue, follows from the interaction of independently motivated, default axiom schemata for inferring particular coherence relations.

Finally, in §4, we summarize our contributions and questions for further research.

*Thanks to Julian Schlöder, Matt Husband, Runyi Yao, Adrian Brasoveanu, Bridget Copley, Fabienne Martin, Roger Schwarzschild, audiences/reviewers at SuB27, XPRAG9, and participants at COCOA. Any errors are ours.

¹[12], [13] and [9] also use 'CIC' to characterize multi-clause examples like *The company fired the manager who was embezzling money*. Here, we use 'CIC' for examples that (at least prima facie) involve a single clause.

2 Experimental support for clause-internal coherence

2.1 Experiment 1: Deverbal adjectives

We hypothesized that attributive deverbal adjectives can trigger clause-internal inferences that parallel those that predicative deverbals trigger at the discourse level. To test this, and to probe for effects of (linear) cause/effect order, we used a 2x2 design crossing INFERENCE DOMAIN {DISCOURSE, CLAUSE} with CAUSE/EFFECT ORDER {CAUSE-EFFECT, EFFECT-CAUSE} for 40 items, as in (2). 42 filler items were balanced for causal inference strength (strong/medium/weak). In an initial Likert-scale study, we found evidence that attributive deverbals can trigger causal coherence-type inferences, but more weakly than their predicative counterparts.²

Building on our initial findings, we hypothesized that, in clause-internal contexts, at least two coherence inferences—a causal inference (4a) and a non-causal inference (4b)—compete for salience, and that the less-salient inference is nonetheless available. In discourse-level contexts, on the other hand, we did not expect the causal inference to have any significant competition, either in salience or general availability. We tested this via a two-stage forced-choice task, under the linking hypothesis that the first stage gauged the relative saliences of the provided interpretations, and that the two stages combined gauged the overall availability of each interpretation. Participants³ (n=48) first chose between a causal and a non-causal interpretation, e.g., the options in (4) for the item in (2). On the same screen, participants then answered the question, *Is the other option also a reasonable description of what happened?*

- (4) a. The child was drenched because she got hit by the big water balloon.
 b. The child was already drenched when she got hit by the big water balloon.

| | Causal | | | Non-causal | | |
|------------------|------------|---------|----------------|------------|---------|---------|
| | 1st Choice | 2nd Ch. | Overall Avail. | 1st Ch. | 2nd Ch. | Overall |
| Disc. Eff.-Cause | 0.91 | 0.07 | 0.98 | 0.09 | 0.33 | 0.42 |
| Disc. C-E | 0.95 | 0.03 | 0.98 | 0.05 | 0.19 | 0.24 |
| Clause. E-C | 0.61 | 0.25 | 0.86 | 0.39 | 0.32 | 0.71 |
| Cl. C-E | 0.61 | 0.24 | 0.85 | 0.39 | 0.33 | 0.72 |

Table 1: Rates of interpretative choices for deverbal adjectives

The proportions of causal and non-causal choices are summarized in Table 1. Data were analyzed in R with maximal Bayesian mixed effects linear regression models [7, 8]. The causal interpretation was a both more frequent first choice (3.89, [2.78, 5.16]) and more available overall (2.65, [1.21, 4.33]) in DISCOURSE conditions, compared to CLAUSE conditions. With respect to the overall availability of the non-causal interpretation, we found significant main effects of both factors: overall availability was higher for CLAUSE conditions than DISCOURSE conditions (-2.54, [-3.16, -1.97]) and higher for EFFECT-CAUSE order than CAUSE-EFFECT (-0.53, [-0.82, -0.24]). There was also a significant interaction: in DISCOURSE conditions, the non-causal interpretation was more available overall in the EFFECT-CAUSE condition than the CAUSE-EFFECT condition (-1.18, [-1.56, -0.82]), but no difference obtained between the CLAUSE conditions.

These results suggest that: (i) attributive deverbal adjectives can reliably trigger inferences that parallel discourse-level coherence inferences, but (ii) attributive deverbals are less likely to

²For full details of this study and a single-stage forced choice study, see [19].

³Participants were native-English-speaking, UK-based Prolific workers. Experiments ran on Prolific [23].

trigger causal coherence inferences than their predicative counterparts. This may be because a non-causal inference can more readily compete with the causal inference in clause-internal contexts, though in discourse this inference is at least available in many of the tested cases.

2.2 Experiment 2: Non-deverbal adjectives

Expt. 1 focused on deverbal adjectives in order to elicit the most robust clause-internal coherence inferences; deverbal adjectives, while not overtly clausal, are derivationally related to verbs and describe events. In Expt. 2 ($n = 60$), we tested whether adjectives lacking these properties also trigger causal inferences when used attributively. The task was the same as in Expt. 1, but the critical adjectives were non-deverbal—e.g., *wet* instead of *drenched* in (2). The results are summarized in Table 2. Data were analyzed in the same way as the Expt. 1 data.

| | Causal | | | Non-causal | | |
|------------------|------------|---------|----------------|------------|---------|---------|
| | 1st Choice | 2nd Ch. | Overall Avail. | 1st Ch. | 2nd Ch. | Overall |
| Disc. Eff.-Cause | 0.86 | 0.04 | 0.90 | 0.14 | 0.31 | 0.45 |
| Disc. C-E | 0.90 | 0.03 | 0.93 | 0.10 | 0.32 | 0.42 |
| Clause. E-C | 0.46 | 0.18 | 0.64 | 0.54 | 0.16 | 0.70 |
| Cl. C-E | 0.44 | 0.20 | 0.64 | 0.57 | 0.13 | 0.70 |

Table 2: Rates of interpretative choices for non-deverbal adjectives

We found that the causal interpretation was a more frequent first choice for DISCOURSE conditions than CLAUSE conditions [3.29, (2.70 3.92)]. There was also a significant interaction: between the DISCOURSE conditions, the causal first choice was more frequent for the CAUSE-EFFECT order (0.65, [0.11, 1.26]), but the CLAUSE conditions did not differ from one another. The causal interpretation was also more available overall for the DISCOURSE conditions (2.55, [1.99, 3.17]). The non-causal interpretation was both a more frequent first choice and more available overall (-1.40, [-1.76, -1.04]) in the CLAUSE conditions than the DISCOURSE conditions.

We also analyzed the combined results of Expts. 1-2, with ADJECTIVE TYPE as an additional factor. We found an effect of ADJECTIVE TYPE such that the causal interpretation was the more frequent first choice (1.16, [0.44, 1.94]) and more available overall (2.03, [1.23, 2.91]) with deverbals; the non-causal interpretation was the more frequent first choice for non-deverbals, and was in fact the first choice over half the time for non-deverbals, but not for deverbals. However, ADJECTIVE TYPE did not have a significant main effect on its overall availability.

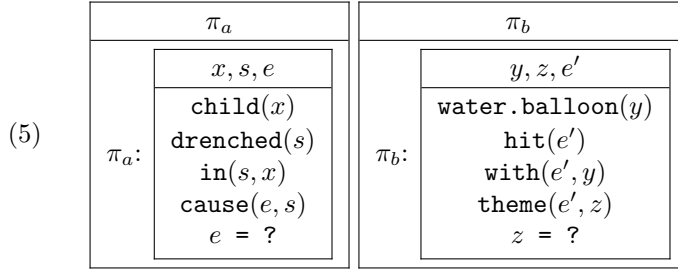
The results of Expt. 2 suggest that non-deverbal adjectives can trigger coherence inferences, even when attributive. However, regardless of context, the causal inferences triggered by non-deverbals seem to be both less salient and less available overall than those triggered by deverbals.

3 Formal Proposal

3.1 Analyzing deverbal adjectives

Discourse coherence. We first consider (5), which illustrates how the discourse-level inference in (2a) is established. According to the asserted content in π_a , there is a child x who is in a drenched state s , which, we propose, was caused by some event e . According to the asserted content in π_b , there is a water balloon y , which was used in a hitting event e' , whose theme is

some individual z . By establishing the coherence relation Elaboration⁴ between π_a and π_b , e is resolved to e' and z is resolved to x .⁵ This resolution entails a causal link between π_a and π_b , although Elaboration is a non-causal relation. The relevant paraphrase is, ‘A drenched state that some child was in was caused by some event; that event is a water-balloon hitting event’.



Let’s now turn to the non-causal interpretation of (2a): ‘A water balloon hit an already drenched child; some other event brought about the drenched state.’ This interpretation follows from establishing Background, which entails that the eventualities described by its arguments overlap in time. For (5), Background ensures that s overlaps e' and that e is bound, but, crucially, not resolved to e' . Assuming that a coherence relation that leaves a discourse referent (dref) unresolved is dispreferred to one that resolves all drefs, our formalization is consistent with the Expt. 1 finding that the causal interpretation is preferred for (2a).⁶

As for (2b), we again have the representational content in (5), but π_a and π_b are interpreted in reverse order. This means that the event causing the drenched state is resolved anaphorically, instead of cataphorically, i.e., the second sentence of (2b) doesn’t expand on the balloon-hitting event, but rather describes its result state. Elaboration is ruled out; we propose that the operative coherence relation here is Continuation⁷. It is non-causal like Elaboration, but supports an anaphoric resolution in which the event causing the drenched state is identified with the balloon-hitting event, thus deriving the causal inference. Finally, while a non-causal inference is not prominent in (2b), it can nevertheless follow from establishing Background, as in (2a).

Clause-internal coherence. We now show how an SDRT analysis of presupposition [4] can be extended to capture (2c) given our hypothesis that attributive deverbals are presupposition triggers. Consider (6). According to the asserted content in π_a , there is an event e of being hit by a water balloon and a child x is e ’s theme. According to the presupposed content in π_b (contributed by the deverbal), there is a drenched state s , which was caused by an event e' , and which holds of an individual z . The presupposed content is related to an underspecified attachment point u by an underspecified coherence relation R . Following our analysis of (2a), we can resolve R to Elaboration or Background, with u being resolved to π_a in either case.⁸ As before, Elaboration leads to a causal interpretation, and Background to a non-causal one.

As for (2d), its representation is the same as (6), with one important distinction: the first argument of R is now u (which is resolved to π_a). Following our analysis of (2b), we can

⁴Establishing Explanation instead would result in a bizarre interpretation: the water-balloon hitting event is the reason for a drenched state being caused by the water-balloon hitting event. However, Explanation is crucial for adjectives that don’t semantically contribute **cause**(e, s), e.g., non-deverbal adjectives.

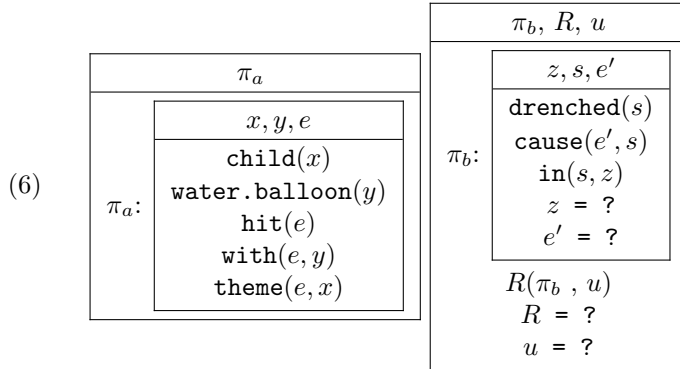
⁵Here we assume that establishing rhetorical relations and resolving the interpretation of a context sensitive expression are correlated tasks (see, e.g., [11, 17, 15, 21]). The semantics of Elaboration ensures that its second argument provides more information about the *same event* described by its first argument (see [5]).

⁶At least one other factor likely underpins this finding: an interpretative default, well-established in experimental psychology [10, 24], to infer a causal link between adjacent eventualities whenever possible.

⁷Establishing Result here would lead to the same bizarre interpretation that was discussed in fn.4.

⁸ π_a is thus either a cataphoric presupposition or a postsupposition (see [6]). We are agnostic about which.

resolve R to either Continuation or Background. As before, Continuation leads to a causal interpretation, while Background leads to a non-causal one.



We are now ready to address two core observations from Expt. 1: (i) causal inferences are more salient and more available in (2a)-(2b) than (2c)-(2d); (ii) for each condition, causal inferences are more salient and available than non-causal inferences. We propose that (i)-(ii) can be captured by assuming two axioms that help adjudicate between possible relations⁹:

- (7) *Resolve drefs*: Establish the relation that produces the least unresolved drefs.
- (8) *Constraint on presuppositions*: If possible, resolve R with Background.

While the discourse-level Elaboration and Continuation respect both axioms, their clause-internal counterparts both violate (8). Yet, despite violating (8), the presuppositional Elaboration and Continuation were still the preferred interpretations in Expt. 1. We can derive this by further assuming that (7) takes precedence over (8).

Our proposal does not presently capture the experimental finding that non-causal inferences are more available overall for (2c) and (2d) than for (2a) and (2b).

3.2 Analyzing non-deverbal adjectives

Let us finally consider non-deverbal adjectives like *happy*. First, we adopt an independently motivated axiom from Schlöder [20, Ch.7]:

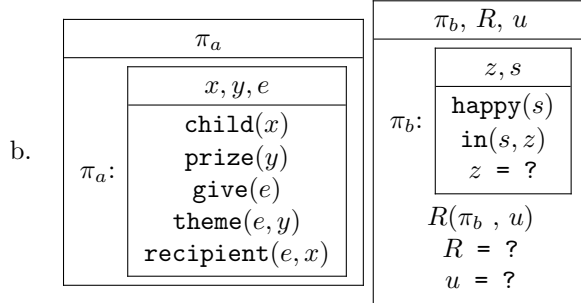
- (9) *Schlöder's causal axiom*: Given a pair of eventuality descriptions α, β :
 - a. if it's possible that the eventuality described by α caused the eventuality described by β , establish $\text{Result}(\alpha, \beta)$.
 - b. if it's possible that the eventuality described by α was caused by the eventuality described by β , establish $\text{Explanation}(\alpha, \beta)$.

We assume that non-deverbal adjectives differ from deverbal adjectives in not semantically contributing $\text{cause}(e, s)$, but that, like deverbal adjectives, they presuppose the described state. The latter is evidenced by (10), which shows the expected projection behavior. Thus, we analyze (11a) as in (11b).

- (10)
 - a. It's not the case that a happy child was given a prize. She was given a new toy.
 - b. A few children at Camp Hope showed up to dinner happy. If a happy child was given a prize, then someone smuggled the prize into the camp.

⁹While (7) is straightforward, we note that (8) is motivated by the intuition that presuppositions are not-at-issue or backgrounded content (see, e.g., [1] and references therein).

(11) a. A happy child was given a prize.



According to the asserted content in π_a , there is a child x who is the recipient of a giving event e , whose theme is a prize y . According to the presupposed content in π_b , there is a state s of being happy that holds of an individual z . This content is related to an attachment point u by a coherence relation R . Unlike in (2c), Elaboration is not a possibility since non-deverbal adjectives don't, by assumption, presuppose **cause**(e, s). However, precisely because they don't presuppose **cause**(e, s), Explanation is possible (recall fn.4), which derives the causal interpretation: the child's happiness is explained by the prize that they were given. This respects Schlöder's axiom, but violates (8). Of course, we can also derive a non-causal interpretation via Background, which would violate Schlöder's axiom but respect (8).¹⁰ Recall that Expt. 2 revealed that the non-causal interpretation (with Background) is the more salient and more available reading for (11a), suggesting that (8) outweighs Schlöder's axiom.¹¹

Let us now consider (12). This example would be analyzed as in (11b), except that the first argument of R would now be u (resolved to π_a). We can derive the non-causal reading by resolving R to Background, and the causal reading by resolving R to Result. Our findings for (12) suggest, as they did for (11a), that (8) outweighs Schlöder's axiom.

(12) A prize was given to a happy child.

Finally, we note that in the discourse versions of (11a)/(12), the axiom in (8) isn't operative. Given Schlöder's axiom, a causal inference is thus most salient via Explanation/Result.

4 Conclusion

We have provided experimental support for the existence of CIC via offline evidence that attributive (non-)deverbal adjectives can trigger the same causal inferences within clauses that their predicative counterparts can trigger across clauses, albeit more weakly. Our forced choice methodology probes whether speakers accept interpretations that are spelled out for them, not which interpretation(s) they might independently infer. Future experiments will investigate what kinds of coherence inferences, if any, speakers draw in the absence of explicit prodding.

To explain the experimental results, we used tools in SDRT, which allowed us to show that causal inferences can be derived in various ways, depending on whether deverbal adjectives are used attributively or predicatively. If the former, they are presupposition triggers and Elaboration/Continuation competes with Background; if the latter, Explanation/Result competes with Background. These different competitions (cash out in terms of interaction between default axioms) correlate with the difference in the relative salience of the causal inferences.

¹⁰The axiom in (7) is not in play here, as non-deverbals do not introduce an event *dref* to resolve.

¹¹For further discussion of Schlöder's axiom, including other constraints that it competes with, see [2].

References

- [1] Márta Abrusán. Presuppositions. In Daniel Altshuler, editor, *Linguistics meets Philosophy*, pages 439–469. Cambridge University Press, 2022.
- [2] Daniel Altshuler. A puzzle about narrative progression and causal reasoning. In Andreas Stokke and Emar Maier, editors, *The Language of Fiction*. Oxford University Press, Oxford, 2021.
- [3] G. E. M. Anscombe. Under a description. *Noûs*, 13:219–233, 1979.
- [4] Nicholas Asher and Alex Lascarides. The semantics and pragmatics of presupposition. *Journal of Semantics*, 15(2):239–299, 1998.
- [5] Nicholas Asher and Alex Lascarides. *Logics of Conversation*. Cambridge University Press, Cambridge, England, 2003.
- [6] Adrian Brasoveanu and Anna Szabolcsi. Presuppositional too, postsuppositional too. 2013.
- [7] Paul-Christian Bürkner. brms: An r package for bayesian multilevel models using stan. *Journal of statistical software*, 80:1–28, 2017.
- [8] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of statistical software*, 76(1), 2017.
- [9] Jonathan Cohen and Andrew Kehler. Conversational eliciture. *Philosophers’ Imprint*, 21:1–26, 2021.
- [10] Arthur C Graesser, Murray Singer, and Tom Trabasso. Constructing inferences during narrative text comprehension. *Psychological review*, 101(3):371, 1994.
- [11] Jerry Hobbs. Coherence and coreference. *Cognitive Science*, 3:67 – 90, 1979.
- [12] Jerry Hobbs. Clause-internal coherence. In Peter Kühnlein, Anton Benz, and Candace L. Sidner, editors, *Constraints in Discourse 2*, page 15–34. John Benjamins, 2010.
- [13] Jet Hoek, Hannah Rohde, Jacqueline Evers-Vermeul, and Ted J.M. Sanders. Expectations from relative clauses: Real-time coherence updates in discourse processing. *Cognition*, 210, 2021.
- [14] Katja Jasinskaja and Elena Karagjosova. Rhetorical relations. In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann, and Thomas Ede Zimmermann, editors, *Companion to Semantics*. Wiley, 2020.
- [15] Elsi Kaiser and Boutaina Cherqaoui. Effects of coherence on anaphor resolution, and vice versa: Evidence from French personal pronouns and anaphoric demonstratives. In A. Holler and K. Suckow, editors, *Empirical Perspectives on Anaphora Resolution*, pages 51–78. 2016.
- [16] Andrew Kehler. Coherence relations. In Robert Truswell, editor, *The Oxford Handbook of Event Structure*, pages 583–604. Oxford University Press, Oxford, 2019.
- [17] Andrew Kehler, Laura Kertz, Hannah Rohde, and Jeffrey L. Elman. Coherence and coreference revisited. *Journal of Semantics*, 25(1):1–44, 2008.
- [18] E.J. Krahmer. *Discourse and presupposition*. PhD thesis, 1995. Pagination: vii, 188.
- [19] Kelsey Sasaki and Daniel Altshuler. Clause-internal coherence: A look at deverbal adjectives. In *Proceedings of Sinn und Bedeutung 27*. Charles University, Prague, to appear.
- [20] Julian J Schlöder. Assertion and Rejection,, 2018. ILLC Dissertation Series, DS-2018-07.
- [21] Una Stojnić. *Context-sensitivity in a Coherent Discourse*. PhD thesis, Rutgers University, 2016.
- [22] Rob Van der Sandt. Presupposition Projection as Anaphora Resolution. *Journal of Semantics*, 9(4):333–377, 11 1992.
- [23] Jeremy Zehr and Florian Schwarz. Penncontroller for internet based experiments (ibex). DOI: <https://doi.org/10.17605/OSF.IO/MD832>, 2018.
- [24] Rolf A Zwaan, Joseph P Magliano, and Arthur C Graesser. Dimensions of situation model construction in narrative comprehension. *Journal of experimental psychology: Learning, memory, and cognition*, 21(2):386, 1995.