# Clause-internal causal inferences: Evidence from nouns[1]

Kelsey SASAKI — *University of Oxford*
Hannah ROHDE — *University of Edinburgh*
Daniel ALTSHULER — *University of Oxford*

**Abstract.** A substantial body of pragmatics research has explored discourse-level coherence inferences — the links between clauses and even larger units that bind a series of utterances into a narrative. Hobbs (2010) proposed that there are also *Clause-internal Coherence* (CIC) inferences, such as that in, 'A jogger was hit by a car,' where one is likely to infer that the car accident occurred while the jogger was jogging. Recently, Sasaki and Altshuler (2022, 2023) and Yao et al. (2024) have provided experimental evidence that CIC inferences can be drawn between verbs and adjectives. We provide further evidence for CIC inferences, this time between verbs and nouns. Our findings from three ratings studies in English suggest that both deverbal nouns like 'jogger' and non-deverbal nouns like 'widow' can give rise to causal CIC inferences. This result is significant because it demonstrates the robustness and pervasiveness of CIC inferences, and raises the question of how proposition-like content may be extracted from a nominal element. We propose a formal analysis of CIC inferences with nouns that adopts the key claim of Pure Event Semantics (Schwarzschild, 2024), namely that nouns always describe eventualities. We synthesize this with a core assumption of Segmented Discourse Representation Theory (Asher & Lascarides, 2003), namely that the arguments of coherence relations are eventuality descriptions. We argue that this provides us with at least two pathways to formally model CIC with nouns.

**Keywords:** clause internal coherence, discourse coherence, experimental pragmatics, SDRT, Pure Event Semantics.

## 1. Introduction

A vital part of interpreting language is the recovery of coherence, i.e., how a series of utterances fit together to form a coherent story. Research in pragmatics has, accordingly, explored discourse-level coherence, which forms the basis for formal theories of discourse structure (see, e.g., Asher and Lascarides, 2003; Asher and Vieu, 2005). However, Hobbs (2010) proposed that coherence inferences also manifest within a clause—between a noun and a verb, for instance. Hobbs provided the following contrast as an illustration: one is likely to infer in (1) that the car accident occurred while the jogger was jogging, while the analogous inference fails to arise in (2); one does not infer that the car accident occurred while the teacher was teaching.

(1) A jogger was hit by a car last night in Marina del Rey.

(2) A teacher was hit by a car last night in Marina del Rey. (adapted from Hobbs, 2010: 16)

Clause-internal coherence (CIC) has received relatively little attention to date, though recent research has motivated its significance. Cohen and Kehler (2021) argue that CIC inferences represent cases of pragmatic enrichment, which they term *eliciture*, that are distinct from familiar types of enrichment such

[1]We would like to thank the audience at S-Side Story (USC Linguistics), especially Elsi Kaiser, Travis Major, and Deniz Rudin for providing feedback on our research. Thanks also to Adrian Brasoveanu, Yue Chen, Jack Duff, Andrew Koontz-Garboden, Matt Husband, Hans Kamp, Paul Pietroski, Julian Schlöder, Roger Schwarzschild, Runyi Yao, and audiences and reviewers at Sinn und Bedeutung 29. Any errors are our own. This research is supported by Leverhulme ECF-2022-604 to Kelsey Sasaki.

as Grice's (1975) implicatures and various forms of local pragmatic strengthening (Levinson, 1987; Recanati, 2010). This work led Sasaki and Altshuler (2022; 2023) to seek experimental evidence that comprehenders actually make CIC inferences. They found that attributive adjectives (e.g., 'the *scared* mouse') can give rise to both causal and non-causal inferences of the kind widely observed at the discourse level. Their findings sharpen the evidence regarding such inferences, showing that CIC inferences are available even between closely related elements such as the verb and the modifier of one of its arguments, thereby extending the experimental evidence beyond sentence-internal effects previously observed only between a matrix clause and a subsequent relative clause (Rohde et al., 2011; Hoek et al., 2021a, b). In the studies reported here, we test whether a CIC inference can arise between a verb and one of its arguments, akin to the Hobbs examples in (1).

To test verb∼adjective inferences, Sasaki and Altshuler manipulated the domain of the inference (discourse vs. clause-internal) and the order of cause and effect, as in (3). Across all four of their studies, they showed that both deverbal (e.g., 'drenched') and non-deverbal (e.g., 'wet') adjectives can give rise to causal inferences at the clause-internal level, though less strongly than at the discourse level. Their findings for cause/effect order in clause-internal contexts were less consistent—in two studies, the causal CIC inference was stronger in effect-cause order (3b) than cause-effect order (3d), but in the other two, no difference obtained.

(3) a. DISCOURSE EFFECT-CAUSE: A child was {drenched / wet}. She got hit by a big water balloon.
b. CLAUSE-INTERNAL EFFECT-CAUSE: A {drenched / wet} child got hit by a big water balloon.
c. DISCOURSE CAUSE-EFFECT: A child got hit by a big water balloon. She was {drenched / wet}.
d. CLAUSE-INTERNAL CAUSE-EFFECT: A big water balloon hit a {drenched / wet} child.
↪ *The child was drenched/wet because she got hit by the big water balloon.*

Following up on this work, Yao et al. (2024) conducted two offline studies probing for causal CIC inferences involving resultative adjectives (e.g., 'broken', 'slippery'). They manipulated the definiteness of the noun phrase (NP) containing the adjective and the cause/effect order. They found that adjectives contained within definite NPs in effect-cause order (4a) gave rise to the strongest causal CIC inferences (relative to (4b)–(4d)).

(4) a. EFFECT-CAUSE, DEFINITE NP: The broken window was struck by a stone.
b. EFFECT-CAUSE, INDEFINITE NP: A broken window was struck by a stone.
c. CAUSE-EFFECT, DEFINITE NP: Bethany struck the broken window with a stone.
d. CAUSE-EFFECT, INDEFINITE NP: Bethany struck a broken window with a stone.
↪ *The window was broken because it was struck by the stone.*

Taken together, Yao et al.'s and Sasaki and Altshuler's findings suggest that causal CIC inferences involving adjectives are generally robust. With respect to cause/effect order, the findings are more mixed: Yao et al.'s findings are broadly consistent with Sasaki and Altshuler's—the causal CIC inference was stronger in effect-cause order, as two of the latter's studies suggested, but only in the definite NP condition ((4a) vs. (4c)). Sasaki and Altshuler only used indefinite NPs, so the two sets of studies are not perfectly comparable. However, between Yao et al.'s findings and Sasaki and Altshuler's, it seems that there is an effect of cause/effect order that is fairly subtle, but consistently points in the same direction when it does appear.

In the current work, we investigated whether nouns can give rise to causal CIC inferences. In §2, we present evidence from three offline studies that both deverbal nouns like 'jogger' and non-deverbal nouns

like 'widow' can give rise to causal CIC interpretations. This result highlights the pervasiveness of such inferences and raises the question of how proposition-like content may be extracted from nominal elements. We did not, however, find a difference in strength between causal CIC inferences involving deverbal nouns and those involving non-deverbal nouns, nor did we find an effect of cause/effect order. This lack of an effect of noun status and cause/effect order contrasts with the findings from Sasaki and Altshuler (2022, 2023) and Yao et al. (2024), which raises questions about why deverbal adjectives may differ from deverbal nouns in supporting stronger causal CIC inferences than their non-deverbal counterparts and whether cause/effect order influences CIC with nouns at all. In §3, we provide a glimpse of what a formal analysis of CIC inferences with nouns may look like and outline its consequences. Our analysis synthesizes two key assumptions from, respectively, Segmented Discourse Representation Theory (SDRT, Asher and Lascarides, 2003) and Pure Event Semantics (Schwarzschild, 2024), namely that arguments of coherence relations are eventuality descriptions and nouns always describe eventualities. In §4, we summarize our contributions and highlight some questions for future research.

## 2. Experimental support for clause-internal coherence with nouns

We hypothesized that nouns, like adjectives, can participate in causal CIC inferences, and that nouns would pattern similarly to the way adjectives did in the studies discussed above. Thus, we had two factors of interest—Cause/Effect Order and Noun Type—based on Sasaki and Altshuler's adjective studies. In a departure from both Sasaki and Altshuler (2022, 2023) and Yao et al. (2024), we also manipulated whether or not a causal inference was expected to obtain at all. In the adjective studies, causal CIC inferences were robustly available in all of the experimental conditions, and therefore ultimately had as a lower baseline only the relatively small number of fillers designed to have a weak (or no) causal CIC inference. This third factor, Expected Causal Inference, gave us a 2x2x2 design. We created the experimental stimuli according to this design, but implemented the actual studies as a series of 2x2s in order to achieve satisfactory statistical power without having to create quite a few more stimuli or recruit an infeasibly large number of participants.

### 2.1. Experiment 1

#### 2.1.1. Design, methods, and predictions

*Design and Materials.* This experiment comprised two sub-experiments.[2] In Experiment 1A, all of the critical nouns were DEVERBAL (e.g., 'winner'); in Experiment 1B, all critical nouns were NON-DEVERBAL (e.g., 'champion'). Both sub-experiments used a 2x2 design crossing Expected Causal Inference {LINK, NO LINK} and Cause/Effect Order {CAUSE-EFFECT, EFFECT-CAUSE} for 40 items, as in (5). Nouns in the LINK condition were selected based on the results of a norming study, the details of which can be found in the Appendix. Nouns in the NO LINK condition were selected based on our own intuitions. 40 one-sentence filler items, which were held constant across sub-experiments, were designed with variable causal inference strengths (weak/medium/strong), different non-verbal CIC triggers (adjectives/nouns), and balanced for cause/effect order.[3]

[2] All studies were run via PCIbex (Zehr and Schwarz, 2018).

[3] Examples of fillers across the three causal inference strengths are below:

(i) a. WEAK: Nalani bumped into a farmer (Graeme)
Question: *How likely do you think it is that Nalani bumped into Graeme because he was a farmer?*

b. MEDIUM: A tipsy partygoer (Peter) was given a drink by Rose.
Question: *How likely do you think it is that Peter was tipsy because Rose gave him a drink?*

c. STRONG: A curious dolphin bumped into a surprised surfer (Jonas).
Question: *How likely do you think it is that Jonas was surprised because the dolphin bumped into him?*

(5) a. LINK EFFECT-CAUSE: Bob congratulated a {winner$_{DEV}$ / champion$_{NDV}$ (Alice).
b. NO LINK EFFECT-CAUSE: Bob congratulated a {reader / bookworm} (Alice).
c. LINK CAUSE-EFFECT: A {winner / champion} (Alice) was congratulated by Bob.
d. NO LINK CAUSE-EFFECT: A {reader / bookworm} (Alice) was congratulated by Bob.
Question: *How likely do you think it is that Bob congratulated Alice because she was a {winner / reader / champion / bookworm}?*

*Participants.* For each sub-experiment, participants were 40 UK-based, native English speakers recruited via Prolific.
*Task.* On a 1–4 scale (*Not at all likely –Extremely likely*), participants responded to a question like that in (5).
*Analysis.* Data were analyzed in R with maximal Bayesian cumulative link mixed effects models using the `brms` package (Bürkner, 2017; Carpenter et al., 2017).
*Predictions.* We predicted a main effect of Expected Causal Inference such that LINK conditions would yield higher causal CIC ratings than NO LINK conditions. On the hypothesis that the order effects observed in the verb∼adjective domain extend to the verb∼noun domain, we also predicted an interaction such that EFFECT-CAUSE is rated higher than CAUSE-EFFECT in the LINK condition, but does not differ in the NO LINK condition. We expected this pattern to either hold across both sub-experiments, or else emerge in Experiment 1A (deverbal nouns) only, because Sasaki and Altshuler (2023) found it for deverbal adjectives, but not non-deverbal adjectives.

### 2.1.2. Results and discussion

The distribution of ratings for experimental stimuli is plotted in Figure 1a. We confirmed that the fillers, which were designed to yield weak, medium, or strong causal inferences, patterned as expected, and that participants used the full rating scale (Fig. 1b).

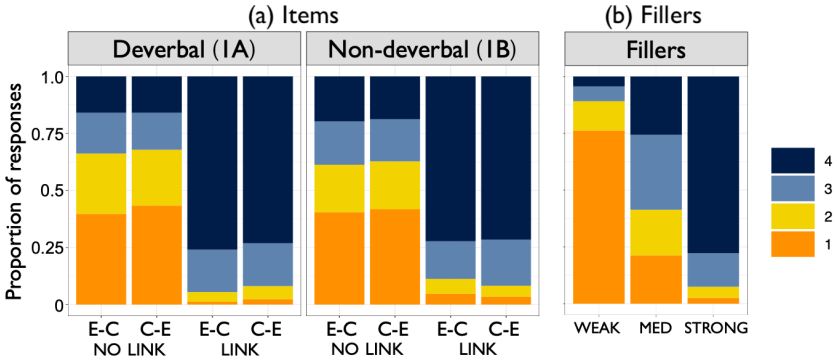


Figure 1. Proportions of ratings for items (a) and fillers (b), Expt. 1.

In both sub-experiments, we found the predicted main effect of Expected Causal Link: causal inferences in the LINK conditions were rated more likely than in the NO LINK conditions with both deverbal nouns ($\hat{\beta}$=4.60, 95% CrI=[3.71, 5.52]) and non-deverbal nouns ($\hat{\beta}$=3.86, CrI=[3.00, 4.75]). Against our predictions, we found no interaction with Cause/Effect Order in either sub-experiment (nor a main effect of Cause/Effect Order).

These results suggest that nouns, like adjectives, can indeed contribute to causal CIC inferences, at least offline. However, the absence of any other effects suggests that they may not pattern similarly to adjectives beyond that. It is possible that this divergence is due in part to the design differences between the adjective studies and the current ones; we leave it to future work to investigate whether Cause/Effect

Order ever affects causal CIC with nouns.[4]

In a post-hoc comparison between the sub-experiments, we did not find an effect of Noun Type. This, too, diverges from the previous findings for adjectives, for which the causal inferences triggered by deverbals (e.g., 'drenched') were stronger than those triggered by non-deverbals (e.g., 'wet'). However, we hesitate to draw a strong parallel between the current results and the adjective findings. The main reason for this is that the adjective studies Sasaki and Altshuler (2023) compared did not use a Likert task, but a two-stage forced-choice task, in which a causal CIC interpretation was pitted against a non-causal CIC interpretation. It is possible that the difference Sasaki and Altshuler observed between deverbal and non-deverbal adjectives only emerges when participants are explicitly presented with multiple interpretations; in the current task, they were only provided with one. Further, the statistical comparisons made here (and by Sasaki and Altshuler) were post-hoc and between-subjects; the absence of a difference between the current sub-experiments could be due to insufficient statistical power. We sought to address this in Experiment 2.

## 2.2. Experiment 2

Neither Experiment 1 nor Sasaki and Altshuler's (2023) studies afforded a within-subjects comparison of deverbal and non-deverbal elements. In this experiment, we treated Noun Type as a within-subjects factor, in place of Cause/Effect Order, which did not have an effect in Experiment 1.[5]

### 2.2.1. Design, methods, and predictions

*Design.* We used a 2x2 design crossing Expected Causal Inference {LINK, NO LINK} with Noun Type {DEVERBAL, NON-DEVERBAL} for 40 items drawn from the Experiment 1 stimuli, as in (6). Items were balanced for cause/effect order. We used the same 40 fillers as in Experiment 1.

(6) a. LINK DEVERBAL: Bob congratulated a winner (Alice).
b. NO LINK DEVERBAL: Bob congratulated a reader (Alice).
c. LINK NON-DEVERBAL: Bob congratulated a champion (Alice).
d. NO LINK NON-DEVERBAL: Bob congratulated a bookworm (Alice).
Question: *How likely do you think it is that Bob congratulated Alice because she was a {winner / reader / champion / bookworm}?*

*Methods.* Participants (N=40) were UK-based, native English speakers recruited via Prolific. The task was the same as in Experiment 1, as was the analysis.
*Predictions.* As in Experiment 1, we predicted a main effect of Expected Causal Inference such that ratings for LINK conditions would be higher than for NO LINK conditions. Under the hypothesis that nouns and adjectives pattern similarly with respect to causal CIC inferences, we also predicted an interaction such that DEVERBAL is rated higher than NON-DEVERBAL in the LINK condition, while there is

[4]Another distinction between the nouns in our studies and the adjectives in previous work is their roles in the causal relation: the nouns were always the cause, while the adjectives were always the effect. It is not clear how or why this might contribute to the different noun and adjective findings. The shift from effect to cause for the nonverbal CIC element in the current studies had a practical motivation: we wanted to compare deverbal and non-deverbal nouns, but struggled to find enough synonymous pairs that could function as the effect. We speculate that this apparent asymmetry may not be accidental, but a detailed discussion of what may underlie it is beyond the scope of this paper.
[5]Further, there was not a consistent effect of Cause/Effect Order across Sasaki and Altshuler's adjective studies. In their Likert task and one-stage forced-choice task, the causal CIC interpretation was, respectively rated higher and chosen more frequently. However, in their two-stage forced-choice tasks, there were no effects of Cause/Effect Order in CIC conditions.

no difference in the NO LINK condition.

### 2.2.2. Results and discussion

The ratings for experimental stimuli are in Figure 2a. We confirmed that participants used the full rating scale and rated the fillers as expected (Fig. 2b).

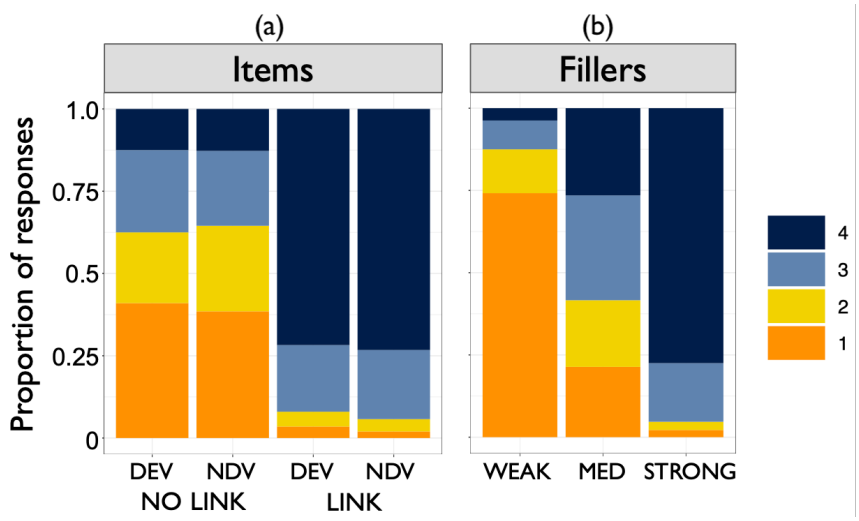


Figure 2. Proportions of ratings for items (a) and fillers (b), Expt. 2.

We found the anticipated main effect of Expected Causal Inference ($\hat{\beta}$=4.48, CrI=[3.63, 5.36])—the causal inference was judged more likely in the LINK conditions in the NO LINK conditions. However, against our predictions, we found no interaction (nor a main effect of Noun Type).

These results provide further evidence that causal CIC inferences can arise between nouns and verbs, at least offline. By virtue of manipulating Noun Type within-subjects, this study provides firmer evidence than Experiment 1 that the deverbal/non-deverbal nature of the noun may not influence final-state causal CIC interpretations. In combination with the lack of Cause/Effect Order effects in Experiment 1, this latter result suggests that there may not be robust parallels between CIC with nouns and CIC with adjectives.

## 3. Towards a formal proposal

The goal of this section is two-fold: (i) to provide a glimpse of what a formal analysis of CIC inferences with nouns can look like and (ii) outline the consequences of such an analysis. To that end, we first adopt a key assumption from SDRT that the arguments of coherence relations are eventuality descriptions (Afantenos et al., 2012). This assumption straightforwardly allows for VPs and clauses to be arguments of coherence relations. Adjectival phrases have also often been treated as eventuality descriptions (see, e.g., Gehrke and McNally, 2009; Wellwood, 2016; Arche et al., 2021). In SDRT, then, adjectives could also be arguments of coherence relations. Sasaki and Altshuler (2022, 2023) show that this is a welcome prediction, opening the door for an SDRT account of CIC with adjectives that they explore in detail.

Our current experimental findings suggest that nouns, including ones that are simple (i.e., those not derived from adjectives or verbs), can also be interpreted as eventuality-describing. However, simple nouns are not canonically treated as eventuality descriptions in formal semantics, so we must make additional assumptions in order to capture the CIC (or indeed discourse-level coherence) inferences that we find with nouns. To that end, we follow Schwarzschild (2024), who proposes a Pure Event Semantics (PES) in which nouns are treated as one-place predicates of eventualities like verbs and adjectives. More

specifically, Schwarzschild treats nouns as predicates of states. For example, both 'winner' and 'villain' in (7) and (8) are predicates of states, namely the states of being a winner and a villain respectively. In PES, proper names (e.g., 'Bob' and 'Anjali' in (7) and (8) respectively) are also predicates of states. This view is more controversial and not something that Schwarzschild (2024) considers in detail.[6] In what follows, we will adopt this view for sake of uniformity, though we note that the CIC inferences that we are after do not depend on this particular treatment of proper names.[7]

(7) A winner was congratulated by Bob.

(8) Anjali chased a villain.

We will also adopt Schwarzschild's terminology—itself following Ramchand (2005)—in calling the entities of whom a state holds *participants*: for instance, the NP 'a winner' in (7) describes a participant who is in a winner state. Moreover, we follow Schwarzschild's proposal—which builds on Schein (2017)—that thematic relations describe states that event participants are in during the course of the event. For example, in (8), which concerns a chasing event, Anjali is in an agent state and the villain is in a patient state. Finally, we assume, again following Schwarzschild, that multiple states may hold of each participant. For example, in (7), the participant that is in a winner state is also in a patient state, and the participant that is in a state of being Bob is also in an agent state.

With these assumptions in place, we can analyze the semantic content of (7) as follows, where $s \circledast s'$ means that $s$ and $s'$ share all participants.

(9) There are two discourse representations $\pi_1$ and $\pi_2$ such that (a)-(c) hold:
   a. $\pi_1$: there is state $s_1$ of being a winner;
   b. $\pi_2$: there is a state $s_2$ of being Bob and an event $e$ of congratulating;
      an agent state of $e$ – call it $s_3$ – such that $s_2 \circledast s_3$;
      a patient state of $e$ – call it $s_4$ – such that $s_1 \circledast s_4$;
   c. there is a relation R such that R($\pi_1$, $\pi_2$).

Note that (9c) crucially incorporates insight from SDRT: since $\pi_1$ and $\pi_2$ are descriptions of eventualities, they must be related by some coherence relation. This follows from the language for incomplete descriptions of Segmented Discourse Representation Structures (SDRSs) called 'the Language $\mathscr{L}_{ulf}$ of Semantic Underspecification'. In addition to expressing underspecified conditions for scope and anaphora, $\mathscr{L}_{ulf}$ also expresses underspecified information about coherence relations: using the higher-order variable ?, '$?(\pi_1, \pi_2, \pi_0)$' is used to express the information that the discourse units $\pi_1$ and $\pi_2$ are connected by some underspecified coherence relation and this resulting connection is part of a complex discourse unit $\pi_0$.

In SDRT, Glue Logic pragmatically enriches the underspecified LFs, including what coherence relation holds between the two discourse units. In the case at hand, two resolutions are possible. One is R = Result, i.e., Result$(\pi_1, \pi_2)$ holds, which would capture the causal interpretation: the winner state causes the congratulating event. The other possibility is to say that Background$(\pi_1, \pi_2)$ holds, which would capture the non-causal interpretation, which we have shown can also arise in this configuration

[6]His main goal is to account for grammatical number, the mass-count distinction, adjectival modification, count adjectives, diminutives, lexical plurals, duals and mass gender.

[7]It is an open question whether proper names can trigger CIC inferences, e.g., that we can understand (i) on a par with (ii).

(i) Usain Bolt broke his ankle.
(ii) The runner broke his ankle.

(critical noun preceding verb). For nouns like 'winner' and 'champion', Result would be the most likely resolution, while Background would be the most likely resolution for nouns like 'reader' and 'bookworm.'

Note that (8) can be analyzed analogously, except that Explanation$(\pi_1, \pi_2)$ rather than Result$(\pi_1, \pi_2)$ would be salient given the effect-cause order (in the active voice). That is, the chasing event is explained by the villain state (rather than the other way around, with the cause-effect order in the passive voice). And, as in (7), (8) could also give rise to Background$(\pi_1, \pi_2)$, capturing the possibility of a non-causal interpretation.

The major challenge for the analysis just outlined is to explain how the compositional semantics, and noun semantics in particular, can generate *discourse representations* (e.g., $\pi_1$ and $\pi_2$ in (9)). Put differently, the proposed analysis commits us to the view the sentences in (7) and (8) are both full-fledged discourses (their contents are stories), and it is not clear how we can generate a propositional sub-part of this discourse from the meaning of a noun.[8]

We do not attempt to address this difficult question here, but rather outline one other analytical path that is available once PES is adopted. In particular, we can pursue the hypothesis below, which differentiates a relation R with respect to whether it is generated within a clause or in a discourse. In the former case (i.e., in the case of CIC), R relates eventualities rather than eventuality descriptions.

(10) **Differentiating discourse coherence from CIC**: Coherence relations that are clause internal are relations between *eventualities*, while coherence relations that are established in a discourse are relations between *eventuality descriptions*.

If one adopts the above hypothesis, they would then not have to answer the difficult question above because nouns, on this hypothesis, simply do not generate discourse representations of eventuality descriptions. They merely relate eventualities in same way that, for instance, an aspectual operator relates eventualities (see, e.g., Moens and Steedman, 1988; and Altshuler et al., 2019: Ch. 5 for an overview). On the hypothesis in (10), the analysis of (7) would proceed as in (11):

(11) There is:
   a. a state $s_1$ of being a winner and a state $s_2$ of being Bob;
   b. an event $e$ of congratulating;
   c. an agent state of $e$ – call it $s_3$ – such that $s_2 \circledast s_3$;
   d. a patient state of $e$ – call it $s_4$ – such that $s_1 \circledast s_4$
   e. a relation R such that R($s_1$, $e$) .

Note the key difference between (9c) and (11e): in the former we relate two eventuality descriptions, while in the latter we relate a state and and event. As in (9c), we say that R in (11e) can be specified as being causal or non-causal by the pragmatics. However, the difficult question for (11e) is where R comes from. Recall that for (9c), R follows from the logic of SDRT since we are relating eventuality descriptions. But in (11), we would need to explain how the R variable—which picks out a relation between an eventuality described by the verb and an eventuality described by the noun—is generated in the logical form. We do not pursue such an explanation here. Rather, we underscore that PES allows one to pursue at least two options for analyzing CIC with nouns, and each raises some interesting, foundational questions about how meanings of nouns compose with the meanings of verbs. We hope to pursue these questions in subsequent research.

[8]Sasaki and Altshuler (2022, 2023) address this question for adjectives by adopting the view that at least some adjectives are presuppositional, and presuppositions warrant rich discourse representations (see, e.g., Van der Sandt, 1992; Asher and Lascarides, 1998). It it not clear to us that this analytical strategy is available in the case of nouns.

## 4. Conclusion

We have provided some initial evidence that nouns can give rise to causal CIC inferences. This aligns with recent findings that adjectives can participate in CIC inferences. We further hypothesized that CIC with nouns would pattern as CIC with adjectives did, but did not find evidence to support this. In particular, neither the linear order of cause and effect nor the nature of the noun influenced the strength of the causal CIC inference.

However, it is not yet clear how robust these inferences are, as the current task explicitly provided the target inference. Our studies were intended only to probe the degree to which participants were willing to endorse the causal inference, not whether they can draw that inference unprompted. In our ongoing work, we have begun testing whether participants construct causal CIC inferences with nouns unprompted and incrementally, using the Maze task (Forster et al., 2009). Promisingly, our early results suggest that they may indeed do so (Sasaki et al., submitted). This converging evidence for CIC with nouns provides additional empirical motivation for our formal proposal, as it suggests that speakers can extract proposition-like content from nouns. This poses a challenge for canonical treatments of nouns, so we have proposed two pathways towards formally modeling CIC with nouns. Both crucially adopt PES (Schwarzschild, 2024), on which nouns are treated as state descriptions.

Many outstanding questions remain about how we can formally model CIC inferences and how they compare to discourse counterparts. We conclude this paper by outlining several observations that warrant further research.

The first observation is that CIC with nominals do not always involve Result or Background. The classic example in (1), repeated below, exemplifies Narration, that is, the jogging is not understood to *cause* the accident but rather *lead to* or *occasion* the accident:

(1) A jogger was hit by a car last night in Marina del Rey. (Hobbs, 2010: 16)

Ultimately, we would want an analysis that is able to capture a wide array of coherence inferences clause-internally, with both nominals and adjectives, comparing them to their discourse counterparts, as below:

(12) There was a jogger in Marina del Rey last night. He was hit by a car.

The second observation is that it is sometimes unclear whether it is a state or an event associated with the noun that coheres with an event described by the verb. Returning to (7) above, one can ask whether it is the state of being a winner that causes the congratulation, or the event that leads to the winning state that causes the congratulation. Note that this question does not arise in (8), where it is more likely that the villain state is understood to be the cause of the chasing event than some event (or events) that led to the villain state.

The third observation comes from the aforementioned experimental findings that CIC with adjectives and CIC with nouns seem to be different. In particular, whether or not an adjective is deverbal appears to play a role in the relative salience of particular coherence relations. However, in our current studies, we did not observe the same distinction between deverbal and non-deverbal nouns. Why should this be the case?

The fourth and final observation is about potential ways of testing whether the hypothesis in (10) is sound, namely, whether the CIC inferences we have discussed are 'true' coherence inferences, on par with cross-clausal coherence inferences. If they are, we would expect them to be able to participate in and affect discourse structure in the same way as cross-clausal coherence. In SDRT, for instance, there is a discourse-structural constraint on which eventuality descriptions that have already been integrated

into the discourse structure are available to participate in additional coherence relations (Asher and Lascarides, 2003).[9] For instance, the first argument of an Explanation should be able to relate to more than one discourse unit, as shown in cross-clausal form in (13a). In this discourse, the first sentence is is the first argument of both an Explanation—in which the second argument is the second sentence—and an Elaboration—in which the second argument is the last sentence. In CIC form, as in (13b), *drenched* should behave the same way.

(13) a. The child was drenched. She got hit by a water balloon. There was even water inside her shoes.
b. The drenched child got hit by a water balloon. There was even water inside her shoes.

One would also expect clause-internal inferences to be sensitive to violations of this discourse-structural constraint. For example, with cross-clausal coherence, the first argument of a Result should not be able to have more than one tail, as shown in (14a); the same should be true of the clause-internal version (14b).

(14) a. A water balloon hit a child. She got drenched. #It struck her right between the shoulders.
b. A water balloon hit a drenched child. ?It struck her right between the shoulders.

An Elaboration between *hit* and *struck* may be available for (14b), but this seems to only be the case if the relationship between *hit* and *drenched* is interpreted as a Background, not a Result.

## References

Afantenos, S., N. Asher, F. Benamara, M. Bras, C. Fabre, L.-M. Ho-Dac, A. Le Draoulec, P. Muller, M.-P. Péry-Woodley, L. Prévot, J. Rebeyrolle, L. Tanguy, M. Vergez-Couret, and L. Vieu (2012). An empirical resource for discovering cognitive principles of discourse organisation: The ANNODIS corpus. In N. Calzolari, K. Choukri, T. Declerck, M. U. Dogan, B. Maegaard, J. Mariani, A. Moreno, J. Odijk, and S. Piperidis. (Eds.), *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, pp. 2727–34. European Language Resources Association.

Altshuler, D., T. Parsons, and R. Schwarzschild (2019). *A Course in Semantics*. MIT Press.

Arche, M. J., A. Fábregas, and R. Marín (2021). On event-denoting deadjectival nominalizations. *The Linguistic Review 38*(2), 191–231.

Asher, N. (2008). Troubles on the right frontier. *Pragmatics and beyond new series 172*, 29.

Asher, N. and A. Lascarides (1998). The semantics and pragmatics of presupposition. *Journal of Semantics 15*(2), 239–299.

Asher, N. and A. Lascarides (2003). *Logics of Conversation*. Cambridge, England: Cambridge University Press.

Asher, N. and L. Vieu (2005). Subordinating and coordinating discourse relations. *Lingua 115*, 591–610.

Bürkner, P.-C. (2017). brms: An r package for bayesian multilevel models using stan. *Journal of statistical software 80*, 1–28.

Carpenter, B., A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell (2017). Stan: A probabilistic programming language. *Journal of statistical software 76*(1).

Cohen, J. and A. Kehler (2021). Conversational eliciture. *Philosophers' Imprint 21*, 1–26.

[9] In SDRT, this is called the Right Frontier Constraint; it is based on the Right Edge Constraint (Polanyi, 1985, 1988). For detailed discussion of the Right Frontier Constraint, we refer the reader to Asher and Lascarides (2003); Asher and Vieu (2005); Holler and Irmen (2007); Asher (2008); Hunter and Thompson (2022), and references therein.

Forster, K. I., C. Guerrera, and L. Elliot (2009). The maze task: Measuring forced incremental sentence processing time. *Behavior research methods 41*, 163–171.

Gehrke, B. and L. McNally (2009). Frequency adjectives and assertions about event types. In S. I. Ed Cormany and D. Lutz (Eds.), *Proceedings of SALT 19*, pp. 180–197. Linguistic Society of America.

Grice, P. (1975). Logic and conversation. In P. Cole and J. L. Morgan (Eds.), *Syntax and Semantics 3*, pp. 325–338. Academic Press.

Hobbs, J. (2010). Clause-internal coherence. In P. Kühnlein, A. Benz, and C. L. Sidner (Eds.), *Constraints in Discourse 2*, pp. 15–34. John Benjamins.

Hoek, J., H. Rohde, J. Evers-Vermeul, and T. J. Sanders (2021a). Expectations from relative clauses: Real-time coherence updates in discourse processing. *Cognition 210*.

Hoek, J., H. Rohde, J. Evers-Vermeul, and T. J. Sanders (2021b). Scolding the child who threw the scissors: Shaping discourse expectations by restricting referents. *Language, Cognition and Neuroscience 36*(3), 382–399.

Holler, A. and L. Irmen (2007). Empirically assessing effects of the right frontier constraint. In *Discourse Anaphora and Anaphor Resolution Colloquium*, pp. 15–27. Springer.

Hunter, J. and K. Thompson (2022). On the role of relations and structure in discourse interpretation. In D. Altshuler (Ed.), *Linguistics meets Philosophy*. Cambridge: Cambridge University Press.

Levinson, S. C. (1987). Implicature explicated? *Behavioral and Brain Sciences 10*, 722–723.

Moens, M. and M. Steedman (1988). Temporal ontology and temporal reference. *Computational Linguistics 2*, 15–29.

Polanyi, L. (1985). A theory of discourse structure and discourse coherence. In W. Eilfort, P. Kroeber, and K. Peterson (Eds.), *Proceedings of the 21st Meeting of the Chicago Linguistic Society*, pp. 306–22. Chicago, IL: Chicago Linguistic Society.

Polanyi, L. (1988). A formal model of the structure of discourse. *Journal of pragmatics 12*(5-6), 601–638.

Ramchand, G. (2005). Post-davidsonianism. *Theoretical Linguistics 31*(3), 359–373.

Recanati, F. (2010). *Truth-Conditional Pragmatics*. Oxford University Press.

Rohde, H., R. Levy, and A. Kehler (2011). Anticipating explanations in relative clause processing. *Cognition 118*, 339–358.

Sasaki, K. and D. Altshuler (2022). Clause-internal coherence as presupposition resolution. In M. Degano, T. Roberts, G. Sbardolini, and M. Schouwstra (Eds.), *Proceedings of the 23rd Amsterdam Colloquium*, pp. 437–443.

Sasaki, K. and D. Altshuler (2023). Clause-internal coherence: A look at deverbal adjectives. In *Proceedings of Sinn und Bedeutung*, Volume 27, pp. 571–588.

Sasaki, K., D. Altshuler, and H. Rohde (submitted). Incremental interpretation of clause-internal coherence: Evidence from the maze. *11th Experimental Pragmatics Conference (XPRAG)*.

Schein, B. (2017). *'And' Conjunction Reduction Redux*. Cambridge: MIT Press.

Schwarzschild, R. (2024). Pure Event Semantics. Forthcoming in Perspectives. Manuscript available at: https://ling.auf.net/lingbuzz/006888.

Van der Sandt, R. (1992, 11). Presupposition Projection as Anaphora Resolution. *Journal of Semantics 9*(4), 333–377.

Wellwood, A. (2016). States and events for s-level gradable adjectives. In J. C. Mary Moroney, Carol-Rose Little and D. Burgdorf (Eds.), *Proceedings of SALT 26*, pp. 166–184.

Yao, R., E. M. Husband, and D. Altshuler (2024). Topichood and temporal interpretation of DPs guide clause-internal, causal coherence. In *Proceedings of Sinn und Bedeutung*, Volume 28, pp. 1034–1048.

Zehr, J. and F. Schwarz (2018). Penncontroller for internet based experiments (ibex). *DOI: https://doi.org/10.17605/OSF. IO/MD832*.

## 5. Appendix

We conducted a norming study (N=80) in order to identify sentences in which a causal CIC inference could robustly be drawn between the verb and the Patient noun—which was to be the critical noun in our main studies—without requiring substantial support from an Agent noun that was not a proper name. We used a 2x2x2 design crossing Cause/Effect Order {CAUSE-EFFECT, EFFECT-CAUSE} with Agent Nominal Type {COMMON, NAME} with Noun Type {DEVERBAL, NON-DEVERBAL} for 48 items, as in (15). We treated Noun Type as a between-subjects factor. 36 list-invariant fillers were designed in the same manner as the main study fillers, and were additionally balanced for whether they contained one or two common nouns. The task was the same as that used in our main studies.

(15) a. CAUSE-EFFECT, COMMON: A {winner / champion} (Bob) was congratulated by a spectator (Alice).
b. CAUSE-EFFECT, NAME: A {winner / champion} (Bob) was congratulated by Alice.
c. EFFECT-CAUSE, COMMON: A spectator (Alice) congratulated a {winner / champion} (Bob).
d. EFFECT-CAUSE, NAME: Alice congratulated a {winner / champion} (Bob).
Question: *How likely do you think it is that Alice congratulated Bob because he was a {winner / champion}?*

We excluded from our main studies all items in which the difference in mean rating between the COMMON and NAME conditions was at least 1 for either the deverbal or the non-deverbal condition. This resulted in six exclusions, as in (16).

(16) a. A choreographer (Pedro) shouted at a {dancer / ballerina} (Natalie).
b. Pedro shouted at a {dancer / ballerina} (Natalie).
Question: *How likely do you think it is that Pedro shouted at Natalie because she was a {dancer / ballerina}?*

We then excluded the two items with the next largest differences between the COMMON and NAME conditions, yielding the 40 items we used in our main studies.