

Clause-internal coherence: A look at deverbal adjectives¹

Kelsey SASAKI — *University of Oxford*

Daniel ALTSHULER — *University of Oxford*

Abstract. Hobbs (2010) introduced the term *Clause-Internal Coherence* (‘CIC’) to describe inferences such as that in, ‘A jogger was hit by a car,’ where the jogging is understood to be implicated in the car-hitting event. Cohen & Kehler (2021) motivate an account of CIC using tools familiar from discourse coherence research. An outstanding question is how to compositionally derive CIC from coherence relations. We propose that CIC can arise as a byproduct of presupposition resolution, couching our analysis in *Segmented Discourse Representation Theory* (Asher & Lascarides 1998) and providing motivation from experimental findings. Our findings suggest: (i) attributive adjectives, both deverbal and non-deverbal, can trigger CIC; (ii) attributive adjectives trigger weaker causal inferences, but stronger non-causal inferences, than their predicative counterparts; (iii) non-deverbal adjectives are weaker causal inference triggers than deverbal adjectives. We argue that attributive adjectives are presupposition triggers, and that they give rise to CIC inferences as a result of presupposition resolution. Thus, CIC with deverbal adjectives arises via Background (non-causal inference) or, depending on word order, Elaboration or Continuation (causal inference). For non-deverbal adjectives, non-causal inferences also arise via Background, but causal inferences arise via Explanation or Result. Finally, we show how some of the interpretative preferences observed in our studies can be modeled as interactions between independently motivated default axioms for choosing between coherence relations. Our research sheds new light on how presupposition relates to anaphora resolution and coherence, while also contributing to recent work on adjectival meaning in discourse.

Keywords: clause internal coherence, discourse coherence, presupposition, anaphora, SDRT, experimental pragmatics

1. Introduction

The sentences below exemplify what Hobbs (2010) terms *Clause-Internal Coherence* (‘CIC’) inferences: instances of the inferences that characterize discourse (i.e., multi-clausal) interpretation that are ‘special’ by virtue of arising from a single clause.²

For example, (1) gives rise to a causal inference: by sticking a knife into her husband, the described subject became a widow.

(1) A widow stuck a knife into her husband. (Anscombe 1979)

¹Thanks to Julian Schlöder, Matt Husband, Runyi Yao, Oliver Bott, Adrian Brasoveanu, Bridget Copley, John Gluckman, Fabienne Martin, Andrew McKenzie, Hannah Rohde, Roger Schwarzschild, Robert Truswell, audiences and reviewers at the 23rd Amsterdam Colloquium, Sinn und Bedeutung 27 and XPRAG 9, and participants at COCOA and Linguistics Colloquy at the University of Kansas. Any errors are our own.

²Hobbs (2010) and Cohen and Kehler (2021) also use ‘CIC’ to describe multi-clausal sentences like those in (i). Here, we only consider instances of CIC that, at least *prima facie*, involve a single clause.

(i) a. The company fired the manager who was embezzling money.
b. The company fired the manager who was hired in 2002.
c. The company fired the manager who has a long history of corporate awards. (adapted by Cohen and Kehler (2021) from an example of Rohde et al. (2011)).

In (2a), it is natural to infer that jogging led to a car accident, but in (2b), we do not make the analogous inference that teaching led to a car accident.

- (2) a. A jogger was hit by a car last night in Marina del Rey.
- b. A teacher was hit by a car last night in Marina del Rey. (Hobbs 2010)

Cohen and Kehler (2021) argue that CIC cannot be accounted for by familiar pragmatic tools including Grice's (1975) implicatures, Bach's (1994) implicatures, and various forms of local pragmatic strengthening proposed by Levinson (1987), Recanati (2010), and others. Instead, they motivate a novel type of enrichment, *eliciture*, that is characterized by its non-local nature, which is familiar from research on the coherence of intersentential discourse.³ Beyond this, clause-internal coherence has received little formal attention; one of many outstanding questions is how to compositionally derive clause *internal* coherence from clause *external* coherence relations. In this paper, we introduce one possible approach to answering this question.

Clause-internal coherence has not yet received much experimental attention either; another aim of this paper is to deepen our empirical understanding of the phenomenon. In §2, we present the results of four offline experiments that probe the strength, salience, and overall availability of the causal inferences triggered by deverbal and non-deverbal adjectives. As shown in (3), we considered deverbal adjectives like *drenched* and non-deverbal adjectives like *wet* in discourse contexts vs. clause-internal contexts, and in cause/effect vs. effect/cause orders.⁴

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

Our findings are consistent with the hypothesis that attributive (i.e., DP/NP-internal) adjectives can trigger clause-internal coherence inferences that correspond to discourse coherence inferences triggered by predicative adjectives. Specifically, the results suggest that attributive adjectives are weaker triggers of causal inferences, but stronger triggers of non-causal inferences, than their predicative counterparts. We observed this pattern for both deverbal and non-deverbal adjectives, suggesting that the trigger of a CIC inference need not necessarily be event-describing or derivationally related to a verb. However, these characteristics do seem to have some effect on coherence inferences more broadly: the causal inferences triggered by non-deverbal adjectives were both less salient and less available overall than those triggered by deverbal adjectives.

Subsequently, in §3, we propose a formal analysis that captures the experimental findings. Specifically, we explore the possibility of extending an analysis of presupposition in *Segmented Discourse Representation Theory* (SDRT, Asher and Lascarides 1998). This account hinges on the argument that CIC can arise as a presupposition, whose trigger is an attributive adjective. Evidence for the presupposition is provided in (4), which shows that *drenched* and *wet* both trigger an inference that projects out of negative, interrogative and suppositional contexts. In

³For recent overviews of this research, see, e.g., Kehler (2019), Jasinskaja and Karagjosova (2020), Altshuler and Truswell (2022: Ch.6).

⁴See Appendix for more sample stimuli.

particular, both adjectives presuppose the described state (of being drenched and wet, respectively), with the deverbal adjective also presupposing that the state is caused by some event.

- (4)
- a. It is not the case that a {drenched/wet} child got hit by a big water balloon. She was pushed into the pool.
 - b. A: Did a {drenched/wet} child get hit by a big water balloon?
B: No, she was pushed into the pool.
 - c. A few children at Camp Hope showed up to dinner {drenched/wet}. If a {drenched/wet} child got hit by a big water balloon, then someone smuggled such balloons into the camp.

The analysis we propose treats presupposition as a species of anaphora resolution (Van der Sandt 1992; Krahmer 1995). We show how CIC inferences follow from the resolution of a coherence relation that binds the presupposed information and an attachment point that allows for projection. We argue that, when distinct inferences are available, the relative salience of each inference 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.

2. Experimental support for clause-internal coherence

2.1. Experiment 1: Rating causal inference strength

We focused first on deverbal adjectives because we hypothesized that, if clause-internal coherence is a robust phenomenon, deverbals would be more likely to give rise to it than non-deverbals. This is because discourse coherence is largely based on relationships between clausal eventuality descriptions. Although deverbal adjectives are not overtly clausal, they are derivationally related to verbs and describe events.

As such, we began our investigation by employing a Likert-scale task to gauge the relative strength of the causal clause-internal inferences triggered by attributive deverbal adjectives, compared to the causal discourse inferences triggered by their predicative counterparts. We focused on causal interpretations because naive speakers' interpretations of them are fairly straightforward to probe (Singer et al. 1992). Based on informal judgments, we hypothesized that (i) attributive deverbal adjectives give rise to causal CIC inferences, and (ii) the strength of these inferences is modulated by the linear order of the cause and effect descriptions.

2.1.1. Design, Methods, and Predictions

Design. We used a 2x2 design crossing Inference Domain {DISCOURSE, CLAUSE} with Cause/Effect Order {CAUSE-EFFECT, EFFECT-CAUSE} for 40 items, as in (3) above. 42 filler items were balanced between discourse and clause inference domain and for causal inference strength (strong/medium/weak). (For more sample stimuli, including fillers, see the Appendix.)

Participants. Participants were 65 UK-based, native English speakers recruited via Prolific. Participants in all the experiments reported here were also Prolific workers in the same demographic categories.

Task. On a 1–4 scale (*Not at all likely–Extremely likely*), participants responded to a question

of the form, *How likely do you think it is that the child was drenched because she got hit by the big water balloon?*

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 Inference Domain such that CLAUSE conditions would receive more 3–4 ratings than 1–2 ratings, but that they would still be rated lower/less likely than DISCOURSE conditions. Following informal judgments, we also predicted an interaction such that CAUSE-EFFECT would be rated lower than EFFECT-CAUSE for CLAUSE conditions, but not DISCOURSE conditions.

2.1.2. Results and Discussion

The distribution of ratings for experimental stimuli is plotted in Figure 1a. We found the predicted main effect of Inference Domain: causal inferences in DISCOURSE conditions were rated more likely than in CLAUSE conditions (2.33, [1.79, 2.89]). We also found an interaction: between the two DISCOURSE conditions, ratings were higher for the CAUSE-EFFECT order than for EFFECT-CAUSE (0.49, [0.04, 0.94]), while the opposite held between the two CLAUSE conditions (−0.72, [−1.10, −0.36]).

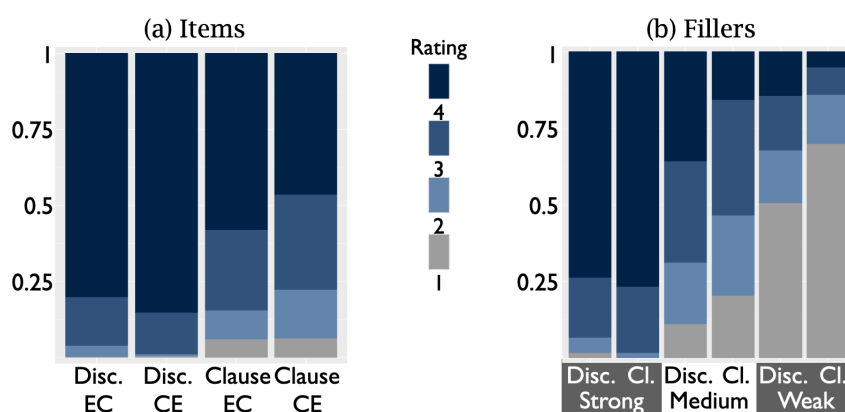


Figure 1: Proportions of ratings for deverbal adjectives, Expt. 1

Ratings were high across the experimental stimuli, but, crucially, this does not seem to have been an experiment-wide pattern. As shown in Figure 1b, the ratings for fillers spanned the full scale, and tracked with the intended strong/medium/weak causal inference strength categories.

The results of this experiment were largely in line with our predictions, thus providing some support for the hypothesis that clause-internal coherence inferences may be reliably and robustly triggered by attributive deverbal adjectives.⁵ One open question is what may be driving the observed Inference Domain effect. We begin to address this in Experiments 2 and 3.

2.2. Experiment 2: One-stage forced choice

We hypothesized that, in Experiment 1, the relative weakness of the causal inferences in CLAUSE contexts compared to DISCOURSE contexts was due to a competing inference in the

⁵We do not undertake a full by-items analysis here, but see §4 for further discussion.

former. Specifically, we posited that this was a non-causal, temporally overlapping inference.⁶

2.2.1. Methods and Predictions

We used the same design and stimuli as in Experiment 1. Participants ($n = 64$) responded to a forced choice question of the form in (5). Our linking assumption was that the task would provide an estimate of the relative salience of each interpretation.

- (5) **A drenched child got hit by a big water balloon.**
 Which is the most accurate description of what happened?
 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.*

We predicted that the non-causal inference (e.g., (5b)) would be chosen more frequently in CLAUSE contexts than DISCOURSE contexts. We also predicted that the non-causal inference would be chosen more frequently for the CLAUSE EFFECT-CAUSE condition than the CLAUSE CAUSE-EFFECT condition. Data were analyzed with maximal Bayesian logistic mixed effects regression models.

2.2.2. Results and Discussion

The rates of non-causal choices are plotted in Figure 2. In line with our predictions, the non-causal interpretation was chosen more frequently for CLAUSE conditions than DISCOURSE conditions (3.77, [2.76, 4.95]). Further, we found an interaction such that, between the two CLAUSE conditions, the non-causal interpretation was chosen more frequently for CAUSE-EFFECT order than for EFFECT-CAUSE order (-0.41 , $[-0.78, -0.07]$) (i.e., the *causal* interpretation was chosen *less* frequently for CE order); no difference obtained between the DISCOURSE conditions.




Figure 2: Rates of non-causal interpretations for deverbal adjectives, Expt. 2

These results suggest that a causal inference and non-causal inference may be in competition in clause-internal contexts. That is, although the causal inference still tends to be the more salient, the non-causal inference can win out more readily than in discourse contexts. However, this method does not show, for a given trial, whether the option that was not chosen was not

⁶In SDRT, this inference is often characterized by the coherence relation Background (Asher et al. 2007). See §3 for more discussion.

chosen because it was an unavailable interpretation, or simply because it was less salient or less plausible, though still available. We sought to address this in Experiment 3.

2.3. Experiment 3: Two-stage forced choice

Based on informal judgments, we hypothesized that the non-causal interpretation would have lower overall availability in discourse contexts than in clause-internal contexts, on top of the lower salience we observed in Experiment 2.

2.3.1. Methods and Predictions

We tested this hypothesis with a two-stage forced choice task, using the same design and stimuli as the previous experiments. The first stage was the same as that in Experiment 2. The second stage was then presented on the same screen. Participants ($n = 48$) responded ‘Yes’ or ‘No’ to the question, *Is the other option also a reasonable description of what happened?*. As in Experiment 2, we made the linking assumption that the first stage response gauges the relative salience of each interpretation. We further assumed that the two stages combined provide an estimate of the overall availability of each interpretation.

We expected that the first choice responses would be consistent with the Experiment 2 findings. We also predicted a main effect of Inference Domain on the total non-causal choices, such that this interpretation was chosen more overall in CLAUSE conditions than DISCOURSE conditions. Data were analyzed with maximal Bayesian mixed effects logistic regression models.

2.3.2. Results & Discussion

The proportions of causal and non-causal choices are plotted in Figure 3. In DISCOURSE conditions, compared to CLAUSE conditions, the causal interpretation was a more frequent first choice (3.89, [2.78, 5.16]) and more available overall (2.65, [1.21, 4.33]). In contrast to Experiments 1 and 2, we found no interaction in the first choice data. There was also no interaction with respect to the overall availability of the causal interpretation.




Figure 3: Rates of interpretative choices for deverbal adjectives, Expt. 3. The lower segment of each bar represents the first choice rate; the upper bar, the second choice rate.

We turn now to the overall availability of the non-causal interpretation.⁷ There were re-

⁷Note that the first choice results are the complement of the causal first choice results.

liable main effects of both factors: overall availability was higher for CLAUSE than DISCOURSE ($-2.54, [-3.16, -1.97]$) and higher for EFFECT-CAUSE order than CAUSE-EFFECT ($-0.53, [-0.82, -0.24]$). There was also an unexpected, but reliable, interaction such that, between the two DISCOURSE conditions, the non-causal reading was more available for the EFFECT-CAUSE order than the CAUSE-EFFECT order ($-1.18, [-1.56, -0.82]$), but no difference obtained between the CLAUSE conditions.

In sum, we found that, for all four conditions, the causal interpretation was both more salient and more available overall than the non-causal interpretation. We also observed higher salience and availability of the causal interpretation in discourse contexts, compared to clause-internal ones; this is in line with Experiments 1 and 2. As for the non-causal interpretation, we found that it was less available overall in discourse contexts than in clause-internal ones. We leave it to future work to unpack which factors, such as temporal iconicity, may be driving the unanticipated effect of Cause/Effect Order in discourse contexts.

2.4. Experiment 4: Two-stage forced choice with non-deverbal adjectives

2.4.1. Design, methods, and predictions

In Experiments 1–3, we focused on deverbal adjectives because, as noted in §1, we hypothesized that deverbals would be more likely to give rise to CIC inferences than non-deverbals. The results suggest that deverbal adjectives can give rise to CIC inferences, but do not indicate the extent to which this effect depends on the adjectives' relation to verbs and event-describing nature. As such, in Experiment 4, we tested whether non-deverbal adjectives can also trigger CIC inferences, particularly causal ones. We hypothesized that non-deverbals would be less likely to do this than deverbals because they are not event-describing, but not that they would be altogether unable to do so, given the psycholinguistic evidence for a causal preference or default in discourse comprehension (Mandler 1986; Zwaan et al. 1995; Briner et al. 2012: a.o.).

The design was the same as in the previous experiments, but the stimuli featured non-deverbal adjectives—for instance, *wet* instead of *drenched* in (3).⁸ The task ($n = 60$) and analysis were the same as in Experiment 3. We predicted the causal interpretation would be more salient and more available overall in DISCOURSE conditions. We also predicted this main effect would be larger here than in Experiment 3.

2.4.2. Results and discussion

The rates of causal and non-causal choices are plotted in Figure 4. With respect to the first choice data, we found that the causal interpretation was chosen more frequently in DISCOURSE conditions than CLAUSE conditions [$3.29, (2.70\ 3.92)$]. This is in line with the results of Experiments 1–3, but the gap was larger in this experiment than the others. For the attributive non-deverbals, the non-causal interpretation was the first choice over half the time; this was not the case among the attributive deverbals. There was also a reliable interaction such that, between

⁸Of the 40 experimental items, 29 had a minimal or near-minimal partner in the deverbal stimuli: minimal partners simply had a non-deverbal adjective in place of the critical deverbal adjective, and near-minimal partners also had a different passive verb (e.g., 'got' instead of 'was') to preserve acceptability. The remaining 11 stimuli were altered more substantially in order to maintain acceptability and causal inference strength (based on our own judgments). Preliminary descriptive analysis does not suggest that these changes had systematic effects on the task results, but passive type is discussed further in §4.

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.




Figure 4: Rates of interpretative choices for non-deverbal adjectives, Expt. 4. The lower segment of each bar represents the first choice rate; the upper bar, the second choice rate.

The causal interpretation was more available overall for the DISCOURSE conditions than CLAUSE conditions (2.55, [1.99, 3.17]). Meanwhile, the non-causal interpretation was more available overall in the CLAUSE conditions than the DISCOURSE conditions (−1.40, [−1.76, −1.04]). These findings track with those of Experiment 3, but unlike in Experiment 3, we observed no effect of Cause/Effect Order, and no interaction.

We also compared deverbals and non-deverbals by analyzing the combined results of Experiments 3 and 4, 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. Thus, our results suggest that non-deverbal adjectives can trigger coherence inferences in both discourse and clause-internal contexts. 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.

2.5. Summary of experimental results

Finally, let us take stock of the main findings of all four experiments, so we can see what our formal analysis ought to account for. First, across all the experiments, the causal interpretation was more dominant in the DISCOURSE conditions compared to the CLAUSE conditions (i.e., rated more likely in Expt. 1 and chosen more frequently in Expts. 2–4). In Expts. 2–4, this means that the first-choice (and only choice, in Expt. 2) non-causal relationship was the opposite. In Expts. 3 and 4, we further found that the non-causal interpretation was more available overall in the CLAUSE conditions than the DISCOURSE conditions. In Expt. 3, we observed an unexpected main effect of Cause/Effect Order on the overall availability of the non-causal interpretation, but a reliable interaction also obtained here.

The interactions we found across the experiments were driven by simple effects of Cause/Effect Order. In this paper, we will not offer an analysis of these effects, and leave it to future experimental work to better determine their robustness. However, as we will see, the observed differences seem to be consistent with one another, rather than multiple distinct or contradictory patterns. We summarize the relevant pairwise comparisons here. First, in line with our

intuitive predictions, CLAUSE EFFECT-CAUSE received higher causal inference ratings (Expt. 1) and more causal interpretation choices (Expt. 2) than CLAUSE CAUSE-EFFECT. Second, in a tidy reversal, DISCOURSE EFFECT-CAUSE received lower causal inference ratings (Expt. 1) and fewer first-choice causal interpretations (Expt. 4) than DISCOURSE CAUSE-EFFECT. In addition, the overall availability of the *non*-causal interpretation was *higher* for DISCOURSE EFFECT-CAUSE than DISCOURSE CAUSE-EFFECT in Expt. 3; this is consistent with the other DISCOURSE effects, if not as directly matching.

In §3, we present our account of how the causal and non-causal interpretations arise across the four experimental conditions. We aim in particular to capture the experimental observations that (i) the causal interpretation was more dominant in discourse contexts than clause-internal contexts, and (ii) within clause-internal contexts, the causal interpretation was dominant for deverbal adjectives, but the *non*-causal interpretation was dominant for non-deverbal adjectives.

3. Formal proposal

We propose a formal account of clause-internal coherence in the SDRT framework (Asher 1993; Lascarides and Asher 1993; Asher and Lascarides 2003). Specifically, we argue that CIC can arise as a presupposition, and can therefore be captured in SDRT via an independently motivated analysis of presupposition (Asher and Lascarides 1998). We further argue that this analysis is compatible with our experimental findings.⁹

3.1. Coherence with deverbal adjectives

3.1.1. Establishing discourse coherence

We first consider how discourse coherence is established in (6).

(6) A child was drenched. She got hit by a big water balloon.

The asserted content of π_a is that there is a child x who is in a drenched state s . On top of that we propose that, when used predicatively, the deverbal adjective may contribute to the assertion that the drenched state was caused by some event e .¹⁰ The asserted content of π_b is that there is a water balloon y , which was used in a hitting event e' , whose theme is some individual z . These contents are represented in (7), using the standard DRT box notation.¹¹

(7)

π_a	
π_a :	x, s, e
	child(x)
	drenched(s)
	in(s, x)
	cause(e, s)
	$e = ?$

π_b	
π_b :	y, z, e'
	water.balloon(y)
	hit(e')
	with(e', y)
	theme(e', z)
	$z = ?$

⁹We do not, however, argue against any alternative accounts, whether in SDRT or another framework. Our aim here is only to demonstrate one viable option for modeling CIC.

¹⁰We are agnostic here about what syntactic structure would lead to such an interpretative possibility, though note that the structure of deverbal adjectives varies between its participial and non-participial ('true-adjectival') uses (cf. *drenched*, which is often participial vs. *punched*, which rarely is). Thanks to John Gluckman (p.c.) for a helpful discussion of this distinction, which we hope to explore in future studies.

¹¹For an introduction to DRT, see, e.g., Kamp et al. 2011.

The causal interpretation results from establishing the coherence relation, Elaboration, between π_a and π_b . The semantics of Elaboration dictates that the second argument provides more information about the same event described by the first argument (Asher and Lascarides 2003). This requirement ensures that e is resolved to e' and z is resolved to x .¹² Elaboration is a non-causal relation in and of itself, but because $\text{cause}(e, s)$ is asserted in π_a , establishing Elaboration here entails a causal link between π_a and π_b . The resulting paraphrase is, ‘A drenched state that some child was in was caused by some event; that event is a water-balloon hitting event.’

It is reasonable to expect that we would have derived this causal interpretation via Explanation, an inherently causal relation. However, because we propose that deverbal adjectives semantically contribute $\text{cause}(e, s)$, establishing Explanation here would result in the bizarre interpretation, ‘The water balloon hitting event is the cause of a drenched state being caused by the water balloon hitting event.’ Thus, Explanation is not the best-fitting relation in this case, but, as we will show in §3.2, it is crucial for deriving a causal interpretation when the lexical semantics does not contribute $\text{cause}(e, s)$.

We turn now to the non-causal interpretation of (6), which can be paraphrased as, ‘A water balloon hit an already drenched child; some other event brought about the drenched state.’ This interpretation follows from establishing the Background coherence relation, which entails that the eventualities described by its arguments overlap in time (Lascarides and Asher 1993). For (6), Background ensures that s overlaps e' and that e is bound. Crucially, however, e is not resolved to e' in this case.

In Experiments 1–3, we found evidence to suggest that, for (6), the causal interpretation was more salient and overall more available than the non-causal interpretation. We propose to capture this with an independently motivated axiom for choosing among potential coherence relations. The axiom, defined in (8), states a preference for resolved *discourse referents* (drefs).

(8) *Resolve drefs*: Establish the relation that produces the least unresolved drefs.

Elaboration, which establishes the causal interpretation, produces no unresolved drefs; the non-causal Background produces one. In accordance with (8), then, we correctly predict the preference for Elaboration. We note, though, that at least one other factor is likely in play here: an interpretative default to infer a causal link between adjacent eventualities whenever possible. We discuss this further in §3.2.1.

Let us now consider the discourse in (9).

(9) A child got hit by a big water balloon. She was drenched.

In this case, we have the same representational content as for (6), but π_b is interpreted before π_a . In this order, the event causing the drenched state is resolved anaphorically, instead of cataphorically. This means that the second sentence of (9) does not provide further information about the water balloon hitting event, but instead describes its result state. Elaboration, therefore, is not available here. Instead, we propose that the Continuation coherence relation is established.¹³ Continuation, like Elaboration, is non-causal, but it entails a causal link in this

¹²Here we assume that establishing coherence relations and resolving the interpretation of a context sensitive expression are correlated tasks (see, e.g., Hobbs 1979; Kehler et al. 2008; Kaiser and Cherqaoui 2016; Stojnić 2016; Stojnić and Altshuler 2021).

¹³Establishing Result here would lead to the same bizarre interpretation that Explanation would for (6).

context because it supports an anaphoric resolution in which the event causing the drenched state is identified with the water balloon hitting event.¹⁴ Finally, the non-causal inference can arise via Background, as it does for (6), regardless of the flipped order.¹⁵ As is the case for (6), this relation is dispreferred for (9) because it results in an unresolved dref, while Continuation does not result in any.

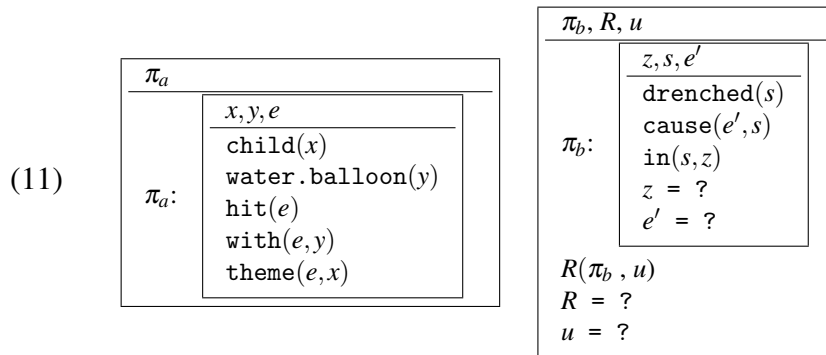
3.1.2. Establishing clause-internal coherence

We hypothesize that attributive deverbal adjectives are presupposition triggers (recall the projectional evidence in (4)). In treating them this way, we can extend Asher and Lascarides' (1998) SDRT analysis of presupposition to derive clause-internal coherence in (10).

- (10) a. A drenched child got hit by a big water balloon.
b. A big water balloon hit a drenched child.

Building on Van der Sandt (1992), Asher and Lascarides propose that the grammar introduces at least two underspecified elements in a presupposition. The first is a coherence relation **R**, which binds the presupposed information; the second is an attachment point **u**, which captures the projectional behavior of presuppositions. Presupposition introduction and resolution are thus part of the usual SDRT discourse update procedure, which ensures that discourse structure and world knowledge can influence the scope of presuppositions.

We first consider the causal interpretation of (10a), represented in (11).



The asserted content in π_a is that there is an event e of being hit by a water balloon and a child x is the theme of e . The attributive deverbal contributes the presupposed content of π_b , namely, that there is a drenched state s , which was caused by an event e' , and which holds of an individual z . Following our analysis of (6), we can resolve R to Elaboration or Background, with u being resolved to π_a in either case.¹⁶ As it does for the discourse-level cases, establishing Elaboration entails a causal link and resolves all the drefs, while establishing Background leads to a non-causal interpretation, but leaves e' unresolved.

¹⁴For discussion of Continuation, see Asher and Lascarides (2003), Altshuler and Truswell (2022: §6.5).

¹⁵Unlike Elaboration, Background is well defined regardless of the ordering of its arguments, e.g., both the discourses below exemplify Background.

(i) a. A man was sitting on a bench. A woman walked over to him.
b. A woman walked over to a man. He was sitting on a bench. (Asher et al. 2007)

¹⁶ π_a is thus either a cataphoric presupposition or a postsupposition. We are agnostic about which. See, e.g., Bott and Solstad (2022) for discussion of the former and Brasoveanu and Szabolcsi (2013) for a discussion of the latter.

The representation of (10b) is nearly the same as that of (10a). The crucial difference is in π_b (see (11)): the first argument of R is u , not π_b . Thus, as it is in the discourse case, Elaboration is ruled out, but we can resolve R to either Continuation or Background. As before, Continuation leads to a causal interpretation, while Background leads to a non-causal one. Assuming the *Resolve drefs* axiom in (8), we can capture the preference for causal CIC inferences over non-causal ones observed in Experiments 2 and 3. Our proposal does not presently capture the experimental finding that non-causal inferences are more available overall in clause-internal contexts than discourse contexts; we leave this for future work.

3.2. Coherence with non-deverbal adjectives

3.2.1. Establishing discourse coherence

As mentioned in §3.1.1, causal coherence must be derived differently when $\text{cause}(e, s)$ is asserted than when it is not. We propose that, while predicative deverbal adjectives semantically contribute $\text{cause}(e, s)$, predicative non-deverbal adjectives do not. The representational content of (12), therefore, is as shown in (13).

- (12) a. A child got wet. She got hit by a big water balloon.
b. A child got hit by a big water balloon. She got wet.

(13)

π_a
π_a :
x, s, e $\text{child}(x)$ $\text{wet}(s)$ $\text{in}(s, x)$

π_b
π_b :
y, z, e' $\text{water.balloon}(y)$ $\text{hit}(e')$ $\text{with}(e', y)$ $\text{theme}(e', z)$ $z = ?$

The absence of $\text{cause}(e, s)$ in π_a clears the path for deriving the causal inferences for (12a) and (12b) via causal coherence relations—Explanation and Result, respectively. The non-causal inference in both cases can be derived by establishing Background. In contrast to the predicative deverbal cases, though, the Backgrounds here do not result in any unresolved drefs. Explanation and Result also resolve all drefs (i.e., z to x), so we must make an additional assumption in order to capture the experimental finding that, with predicative non-deverbal adjectives, the causal interpretation is more salient and more available than the non-causal interpretation.

This brings us to the causal default we gestured to in §3.1.1, namely, to infer a causal relationship between adjacent eventualities whenever possible. This seems to be a robust default, as established by much work in experimental psychology and psycholinguistics (Graesser et al. 1994; Zwaan et al. 1995: a.o.). We port this insight into our SDRT analysis by adopting Schlöder’s axiom (Schlöder 2018: Ch.7):

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

This allows us to formally model the preference for Explanation/Result interpretations.

3.2.2. Establishing clause-internal coherence

Finally, let us consider attributive non-deverbal adjectives. We assume, as we do for predicative non-deverbals, that attributive non-deverbals do not contribute $\text{cause}(e, s)$. However, we assume that attributive non-deverbals are like their deverbal counterparts in presupposing the described state (recall (4)). Thus, we propose to represent (15a) as in (16).

- (15) a. A wet child got hit by a big water balloon.
b. A big water balloon hit a wet child.

(16)

π_a

π_a :

x, y, e
$\text{child}(x)$ $\text{water.balloon}(y)$ $\text{hit}(e)$ $\text{by}(e, y)$ $\text{theme}(e, x)$

π_b, R, u

π_b :

z, s
$\text{wet}(s)$ $\text{in}(s, z)$ $z = ?$

$R(\pi_b, u)$
 $R = ?$
 $u = ?$

The asserted content in π_a is that there is a child x who is the theme of a water balloon hitting event e . The presupposed content in π_b is that there is a state s of being wet that holds of an individual z . In parallel to our proposal for predicative non-deverbals, Elaboration is not a possible value of R here because there is no presupposed $\text{cause}(e, s)$, but its absence makes Explanation possible, without the bizarre interpretation that occurs with deverbals. The non-causal interpretation can be derived via Background.

The content of (15b) is almost identical to that in (16). The only difference is that, in the representation of π_b , the first argument of R is u . Thus, the causal reading is derived by resolving R to Result, and the non-causal reading arises in the now familiar way, via Background.

Finally, we consider how to bring our formal proposal and our experimental findings into concord. The Explanation and Result resolutions respect Schlöder's axiom, while Background violates it.¹⁷ This would predict that the causal interpretations win out, but recall that, in Experiment 4, we found that the *non-causal* interpretation was the more salient and more available reading for (15). Thus, we propose one more axiom, as defined in (17).

- (17) *Constraint on presuppositions*: If possible, resolve R with Background.¹⁸

We propose that this axiom tends to outweigh Schlöder's axiom.¹⁹ This yields a formal outcome that aligns with the experimental observations of Experiment 4. However, attributive deverbal adjectives are also presupposition triggers, so the axiom in (17) is also active when they are involved. We must therefore ensure that we can derive both the causal preference for attributive deverbals and the non-causal preference for attributive non-deverbals. The latter arises from our proposal that the *Constraint on presuppositions* axiom tends to outweigh Schlöder's axiom. To capture the former, then, we propose that the *Resolve drefs* axiom tends to outweigh the *Constraint on presuppositions*. The three axioms tend to be weighted as follows:

¹⁷The *Resolve drefs* axiom in (8) is not in play, as none of these relations produce unresolved drefs.

¹⁸This axiom is motivated by the independently-attested intuition that presuppositions are not-at-issue or backgrounded content (see, e.g., Abrusán (2022) and references therein).

¹⁹For more discussion of Schlöder's axiom, including other constraints that it competes with, see Altshuler (2021).

(18) *Resolve drefs* > *Constraint on presuppositions* > *Schlöder's axiom*.

This is consistent with our experimental findings that the non-causal (Background) interpretation tends to lose out to the causal (Elaboration) interpretation for attributive deverbals, and that the opposite occurs for attributive non-deverbals.

4. Conclusion

We have offered some experimental support for the existence of clause-internal coherence. In our set of offline studies, we found evidence to suggest that attributive (non-)deverbal adjectives can trigger the same causal inferences within clauses that their predicative counterparts can trigger across clauses, albeit more weakly. We also found some support for the hypothesis that a non-causal inference is a more salient competitor to the causal inference when an attributive adjective is involved, compared to a predicative adjective. With these findings, we hope to have contributed to recent work on adjectival meaning in discourse (e.g., Kaiser and Wang 2021). However, we note that our findings do not indicate that attributive adjectives necessarily give rise to CIC inferences, even in the contexts we tested. This is because, in the present experiments, we probed whether speakers accept an interpretation that is spelled out for them (Expt. 1), or which interpretation they preferred, with no option to reject both (Expts. 2–4). That is, we do not yet know which interpretation(s) speakers might independently infer, nor if their preferred interpretation involves a coherence inference. Future experiments will investigate what kinds of coherence inferences, if any, speakers draw in the absence of explicit prompting.

Among our other lines of our ongoing inquiry, we will investigate whether CIC inferences are modulated by the *be*-passive/*get*-passive/*become* distinction. We used both in the current studies to maximize the naturalness of individual stimuli (as shown in the Appendix), but did not explicitly manipulate or control for this factor. We are also conducting by-items analysis of the current studies as a means of informing future investigation of the role lexical semantics may play in CIC. Our descriptive analysis of Expt. 1, for instance, suggests at least one pattern of interest for future systematic testing: for eight of the 40 stimuli, the causal inference strength ratings were nearly identical across the four conditions. The critical adjectives in all eight were derived from psych predicates, including *stunned*, *frustrated*, and *relieved*. However, other stimuli featuring psych adjectives (e.g., *scared*, *elated*, and *annoyed*) patterned similarly to the experiment-wide results.

We have also proposed a formal approach to clause-internal coherence, crucially arguing that attributive adjectives are presupposition triggers. Thus, we were able to use existing tools in SDRT to derive both causal and non-causal CIC inferences, in addition to their discourse counterparts. Then, we showed that some of the robust interpretative preferences observed in the experiments can be modeled as interactions between three independently motivated default axioms for choosing between possible relations. In future work, we will seek to provide a similar account of another seemingly robust finding, namely, the higher overall availability of the non-causal interpretation in clause-internal contexts vs. discourse contexts (Expts. 3–4). Through this and future experimental work, we will continue to test the viability of this approach to modeling clause-internal coherence.

References

- Abrusán, M. (2022). Presuppositions. In D. Altshuler (Ed.), *Linguistics Meets Philosophy*, pp. 439–469. Cambridge University Press.
- Altshuler, D. (2021). A puzzle about narrative progression and causal reasoning. In A. Stokke and E. Maier (Eds.), *The Language of Fiction*. Oxford: Oxford University Press.
- Altshuler, D. and R. Truswell (2022). *The Syntax-Discourse Interface. Extraction from Coordinate Structures*. Oxford: Oxford University Press.
- Anscombe, G. E. M. (1979). Under a description. *Noûs* 13, 219–233.
- Asher, N. (1993). *Reference to Abstract Objects in Discourse*. Dordrecht: Springer.
- 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., L. Prévot, and L. Vieu (2007). Setting the background in discourse. *Discours* 1, 1–30.
- Bach, K. (1994). Conversational implicature. *Mind and Language* 9, 124–162.
- Bott, O. and T. Solstad (2022). Cataphoric presupposition verification and symmetric filtering: The case of occasion verbs. In M. Degano, T. Roberts, G. Sbardolini, and M. Schouwstra (Eds.), *Proceedings of the 23rd Amsterdam Colloquium*, pp. 397–404.
- Brasoveanu, A. and A. Szabolcsi (2013). Presuppositional *Too*, Postsuppositional *Too*. In M. Aloni, M. Franke, and F. Roelofsen (Eds.), *The Dynamic, Inquisitive, and Visionary Life of Φ , $?\Phi$, and $\diamond\Phi$. A Festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman*, pp. 55–64.
- Briner, S. W., S. Virtue, and C. A. Kurby (2012). Processing causality in narrative events: Temporal order matters. *Discourse Processes* 49(1), 61–77.
- 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), 1–32.
- Cohen, J. and A. Kehler (2021). Conversational eliciture. *Philosophers' Imprint* 21, 1–26.
- Graesser, A. C., M. Singer, and T. Trabasso (1994). Constructing inferences during narrative text comprehension. *Psychological Review* 101(3), 371.
- 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. (1979). Coherence and coreference. *Cognitive Science* 3, 67 – 90.
- 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.
- Jasinskaja, K. and E. Karagjosova (2020). Rhetorical relations. In D. Gutzmann, L. Matthewson, C. Meier, H. Rullmann, and T. E. Zimmermann (Eds.), *The Wiley Blackwell Companion to Semantics*, Volume IV, pp. 2645–2673. John Wiley & Sons.
- Kaiser, E. and B. Cherqaoui (2016). Effects of coherence on anaphor resolution, and vice versa: Evidence from French personal pronouns and anaphoric demonstratives. In A. Holler and K. Suckow (Eds.), *Empirical Perspectives on Anaphora Resolution*, pp. 51–78. De Gruyter.
- Kaiser, E. and C. Wang (2021). Packaging information as fact versus opinion: Consequences

- of the (information-)structural position of subjective adjectives. *Discourse Processes* 58, 617–641.
- Kamp, H., J. van Genabith, and U. Reyle (2011). Discourse Representation Theory. In D. Gabbay and G. F. (Eds.), *Handbook of Philosophical Logic*, pp. 125–394. Dordrecht: Springer.
- Kehler, A. (2019). Coherence relations. In R. Truswell (Ed.), *The Oxford Handbook of Event Structure*, pp. 583–604. Oxford: Oxford University Press.
- Kehler, A., L. Kertz, H. Rohde, and J. L. Elman (2008). Coherence and coreference revisited. *Journal of Semantics* 25(1), 1–44.
- Krahmer, E. (1995). *Discourse and Presupposition*. Ph. D. thesis, University of Tilburg.
- Lascarides, A. and N. Asher (1993). Temporal interpretation, discourse relations and common-sense entailment. *Linguistics and Philosophy* 16(5), 437–493.
- Levinson, S. C. (1987). Implicature explicated? *Behavioral and Brain Sciences* 10, 722–723.
- Mandler, J. M. (1986). On the comprehension of temporal order. *Language and Cognitive Processes* 1(4), 309–320.
- 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.
- Schlöder, J. J. (2018). *Assertion and Rejection*. Ph. D. thesis, University of Amsterdam.
- Singer, M., M. Halldorson, J. Lear, and P. Andrusiak (1992). Validation of causal bridging inferences in discourse understanding. *Journal of Memory and Language* 31(4), 507–524.
- Stojnić, U. (2016). *Context-sensitivity in a Coherent Discourse*. Ph. D. thesis, Rutgers University.
- Stojnić, U. and D. Altshuler (2021). Formal properties of *now* revisited. *Semantics & Pragmatics* 14, 1–44.
- Van der Sandt, R. (1992). Presupposition Projection as Anaphora Resolution. *Journal of Semantics* 9(4), 333–377.
- Zwaan, R., J. Magliano, and A. Graesser (1995). Dimensions of situation model construction in narrative comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 21(2), 386–397.

5. Appendix: Sample stimuli

Deverbals (Expts. 1–3)

- (19)
- a. Disc. EC: A tennis pro was stunned. He was beaten by an amateur player.
 - b. Disc. CE: An amateur player beat a tennis pro. The tennis pro was stunned.
 - c. Clause EC: A stunned tennis pro was beaten by an amateur player.
 - d. Clause CE: An amateur player beat a stunned tennis pro.
- (20)
- a. Disc. EC: A barista got overwhelmed. He was bombarded with orders by an afternoon rush of customers.
 - b. Disc. CE: An afternoon rush of customers bombarded a barista with orders. The barista got overwhelmed.
 - c. Clause EC: An overwhelmed barista was bombarded with orders by an afternoon rush of customers.
 - d. Clause CE: An afternoon rush of customers bombarded an overwhelmed barista with orders.

Clause-internal coherence: A look at deverbal adjectives

- (21) a. Disc. EC: An editor was impressed. She was shown a great novel by a young author.
b. Disc. CE: A young author showed a great novel to an editor. The editor was impressed.
c. Clause EC: An impressed editor was shown a great novel by a young author.
d. Clause CE: A young author showed a great novel to an impressed editor.
- (22) a. Disc. EC: A cat got scared. It got chased by a brown dog.
b. Disc. CE: A brown dog chased a cat. The cat got scared.
c. Clause EC: A scared cat got chased by a brown dog.
d. Clause CE: A brown dog chased a scared cat.
- (23) a. Disc. EC: A bicyclist got annoyed. She was cut off by a truck driver.
b. Disc. CE: A truck driver cut off a bicyclist. The bicyclist got annoyed.
c. Clause EC: An annoyed bicyclist was cut off by a truck driver.
d. Clause CE: A truck driver cut off an annoyed bicyclist.
- (24) a. Disc. EC: A river got polluted. It got toxic waste dumped into it by a shady oil company.
b. Disc. CE: A shady oil company dumped toxic waste into a river. The river got polluted.
c. Clause EC: A polluted river got toxic waste dumped into it by a shady oil company.
d. Clause CE: A shady oil company dumped toxic waste into a polluted river.
- (25) a. Disc. EC: A hiker was relieved. He was found by a patrolling park ranger.
b. Disc. CE: A patrolling park ranger found a hiker. The hiker was relieved.
c. Clause EC: A relieved hiker was found by a patrolling park ranger.
d. Clause CE: A patrolling park ranger found a relieved hiker.

Non-deverbals (Expt. 4)

- (26) a. Disc. EC: A tennis pro was speechless. He was beaten by an amateur player.
b. Disc. CE: An amateur player beat a tennis pro. The tennis pro was speechless.
c. Cl. EC: A speechless tennis pro was beaten by an amateur player.
d. Cl. CE: An amateur player beat a speechless tennis pro.
- (27) a. Disc. EC: A barista was distraught. He was bombarded with orders by an afternoon rush of customers.
b. Disc. CE: An afternoon rush of customers bombarded a barista with orders. The barista was distraught.
c. Cl. EC: A distraught barista was bombarded with orders by an afternoon rush of customers.
d. Cl. CE: An afternoon rush of customers bombarded a distraught barista with orders.
- (28) a. Disc. EC: An editor was enthusiastic. She was shown a great novel by a young author.
b. Disc. CE: A young author showed a great novel to an editor. The editor was

- enthusiastic.
- c. Cl. EC: An enthusiastic editor was shown a great novel by a young author.
 - d. Cl. CE: A young author showed a great novel to an enthusiastic editor.
- (29)
- a. Disc. EC: A cat was nervous. It got chased by a brown dog.
 - b. Disc. CE: A brown dog chased a cat. The cat was nervous.
 - c. Cl. EC: A nervous cat got chased by a brown dog.
 - d. Cl. CE: A brown dog chased a nervous cat.
- (30)
- a. Disc. EC: A bicyclist got furious. She was cut off by a truck driver.
 - b. Disc. CE: A truck driver cut off a bicyclist. The bicyclist got furious.
 - c. Cl. EC: A furious bicyclist was cut off by a truck driver.
 - d. Cl. CE: A truck driver cut off a furious bicyclist.
- (31)
- a. Disc. EC: A river was filthy. It got toxic waste dumped into it by a shady oil company.
 - b. Disc. CE: A shady oil company dumped toxic waste into a river. The river got polluted.
 - c. Cl. EC: A filthy river got toxic waste dumped into it by a shady oil company.
 - d. Cl. CE: A shady oil company dumped toxic waste into a filthy river.
- (32)
- a. Disc. EC: A chocolate bar got gooey. It was left in a hot car by a little kid.
 - b. Disc. CE: A little kid left a chocolate bar in a hot car. It got gooey.
 - c. Cl. EC: A gooey chocolate bar was left in a hot car by a little kid.
 - d. Cl. CE: A little kid left a gooey chocolate bar in a hot car.

Fillers (Only Expt. 1 rating prompt shown.)

- (33) Strong causal link:
- a. After an intense training session, the figure skater was very sore. | How likely do you think it is that the skater was sore because the training session was intense?
 - b. Jake's refrigerator was smelly. It had a rotting onion in it. | How likely do you think it is that the refrigerator was smelly because there was a rotting onion in it?
- (34) Medium causal link:
- a. Martin adopted a dog, who was quite elderly. | How likely do you think it is that Martin adopted the dog because it was elderly?
 - b. A little boy pinched his sister on the arm. She poked him in the stomach. | How likely do you think it is that the boy's sister poked him because he pinched her?
- (35) Weak causal link:
- a. A shy bartender saw a tabby cat in her back garden. | How likely do you think it is that the bartender saw the tabby cat because she was shy?
 - b. Celia was a keen baker. She also loved to knit. | How likely do you think it is that Celia was a keen baker because she loved to knit?