Reanalysis in discourse comprehension: Evidence from reading times

John Duff
Saarland University
Saarbrücken, Germany
jduff@lst.uni-saarland.de

Daniel Altshuler
University of Oxford
Oxford, UK
daniel.altshuler@ling-phil.ox.ac.uk

Abstract

This paper explores the evidence for 'discourse reanalysis,' the hypothetical case where an initial interpretation of a sequence of discourse units must be revised due to further context. If discourse reanalysis does occur, our representations may need to permit selective non-monotonic update of discourse-pragmatic meaning like coreference and coherence (Lascarides and Asher 1993; Haug 2014). But a simpler alternative looms: fully-underspecified representations, which could maintain all grammatical discourse interpretations, without any intermediate selection or ranking to later be reanalyzed. To see which approach best models the representations used in actual incremental comprehension, we turn to psycholinguistic evidence. Discourse reanalysis should be associated with momentary difficulty akin to 'garden path' effects in syntactic processing. But previous studies looking for this difficulty in discourse processing have been equivocal. Here, we investigate a case missing from those previous experiments, a joint ambiguity of coreference and coherence, for which we show robust initial preferences. Then, in a self-paced reading experiment, we observe evidence for difficulty when a preferred interpretation must be abandoned. We take this to solidify the empirical basis for discourse reanalysis as a consequence of incremental composition at the discourse level, and discuss two ways this might be accomplished, depending on one's approach to the interface between logical representation and real-time comprehension.

1 Introduction

In some sequences of discourse units $[S_a \dots S_n, S_{n+1}]$, the natural interpretation that would be given for $[S_a \dots S_n]$ alone is not part of the natural interpretation of the completed discourse. For instance, given only S_1 and S_2 of (1), the discourse would seem to describe a series of parallel events, where *him* in S_2 picks out George. But given the continuation in S_3, *him* must instead pick out Bill, and S_1 and S_2 would now seem to describe an action and a reaction.[1]

(1) [Bill pinched George]$_{S_1}$ [and Ava elbowed him.]$_{S_2}$ [Bill didn't elbow her in return,]$_{S_3}$...

In dynamic logics where the first representation of S_1 through S_2 is a step along the construction of the final meaning of the discourse, that representation must somehow be revised in light of S_3. We will say that such representations would have to undergo 'discourse reanalysis.'

As noted by Haug (2014), discourse reanalysis in sequences like (1) poses problems for logics that don't permit selective non-monotonic update of coreference. For instance, in Discourse Representation Theory (DRT; Kamp and Reyle 1993), once a pronoun *him* has been represented as coreferent with its antecedent, the downstream semantics cannot distinguish propositions involving the uncertain referent of *him* from those involving the antecedent. This observation

[1] A reviewer points out that intonation associated with focus structure could play a role in marking this interpretation already at S_2, and wonders whether, as a result, this ambiguity may be somewhat artificial. We are less worried. Because broad focus-marking for S_1 and S_2 permits both interpretations, and given that intonational disambiguation is not especially common in natural production (Allbritton, McKoon, and Ratcliff 1996; Syrett, Simon, and Nisula 2012), we suspect that such ambiguities could arise commonly even in speech. However, we acknowledge that assigning focus structure in silent reading is its own process which certainly interacts with other aspects of interpretive selection (e.g. Carlson et al. 2009), and its role here should be further examined.

motivated Haug (2014) to develop a partial compositional variant of DRT (PCDRT), where dynamic meaning is separated into a monotonically-updated semantic representation, which underspecifies for coreference, and a flexibly-edited set of pragmatic enrichments, through which coreference is ultimately resolved.[2]

To be sure, discourse reanalysis in (1) may involve more than just updating coreference. If the initial interpretation of $[S_1,\ S_2]$ also includes the inference of a coherence relation between the two segments, this relation must also be revised, from Parallel to Result (in the sense of Asher and Lascarides 2003). We follow the standard view that coherence is intricately linked to coreference resolution (e.g. Hobbs 1979; Kehler et al. 2008). On this view, the initial resolution of *him* as George is a consequence of establishing a Parallel relation. But a different coherence relation is preferred once *him* is revised to refer to Bill.[3] Lascarides and Asher (1993) and Asher and Lascarides (2003) use similar examples to argue that coherence too must permit non-monotonic update; to accomplish this in their Segmented Discourse Representation Theory (SDRT), they adopt a glue logic that determines coherence relations through defeasible entailment.

One objection to theories capable of selective non-monotonic update is that a simpler alternative looms: persistent underspecification. Discourse reanalysis could be avoided wholesale if the dynamic representation of the discourse was limited to only an underspecified representation identifying all the interpretations consistent with the monotonic content (cf. common approaches to quantifier scope, see Egg 2010). This would leave open, e.g., uncertain coherence and coreference, only applying defeasible heuristics at the end of the discourse. In such an approach, there is no need for the logic to permit any step-wise non-monotonic update.

So, do we need non-monotonicity here, or would persistent underspecification suffice?[4] If our goal is to model the representations used in actual incremental comprehension, we can turn to behavioral observations from psycholinguistic experiments. True discourse reanalysis would require recognition of an inconsistency, followed by the identification and construction of a suitable revised interpretation; as a whole this should incur momentary difficulty akin to 'garden path' effects in the processing of lexical and syntactic ambiguity (Frazier and Rayner 1982; Duffy, Morris, and Rayner 1988).[5] Readers indeed exhibit slower and more effortful reading when there are inconsistencies in the monotonic content of a text (Albrecht and O'Brien 1993). Yet, to date, reading studies have generally failed to observe clear difficulty for cases of hypothetical discourse reanalysis, suggesting that readers may not experience inconsistency, or launch any subsequent process of costly revision for, e.g., ambiguous pronouns (Stewart, Holler, and Kidd 2007), temporal ordering (Dickey 2001; Sasaki 2021), and coherence-related causal inferences (Mak and Sanders 2013; Duff 2023). This could be explained if readers persistently underspecify interpretations at the level of discourse, as a rule.

In the present work, we push this body of evidence further, looking for the first time (we believe) at the incremental comprehension of discourses with joint ambiguities of coherence and

[2]For an application of this approach to fictional narrative, which often seems to provoke discourse reanalysis, see Altshuler and Haug (2017), Altshuler and Haug (forthcoming), and Altshuler and Kim (2024). See also Kameyama (1996) and Jaspars and Kameyama (1998) for another non-monotonic approach to coreference.

[3]As a reviewer points out, in SDRT it is possible to represent multiple relations, so that a final representation here would, in addition to Result, possibly maintain a weak Parallel relation (weak in that it relies on mere semantic similarity without argument identity). Such an analysis could be built without non-monotonic edits to coherence, depending on one's willingness to equate this weaker Parallel to its stronger predecessor. Nevertheless, this derivation would still require non-monotonic edits to coreference.

[4]Asher and Lascarides (2003) highlight that SDRT is flexible about when pragmatic heuristics are compiled to deliver a preferred interpretation. It is in principle possible in SDRT to provide and revise "best guesses" at every discourse unit, or to retain persistent underspecification and avoid discourse reanalysis. Given this, we could refine our question: Is there any evidence that our representations need to permit the former?

[5]Out of an abundance of caution, we will avoid calling cases of discourse reanalysis 'garden paths' *per se*, as we make no claim that they instantiate the same high degree of difficulty. See the papers in Fodor and Ferreira (1998) for careful discussion of the individual components of that difficulty and how they might vary.

Table 1: An example stimulus set from Experiment 1, following "The kids were misbehaving at a fancy dinner."

Pronoun	**Context**	**Text**
Ambiguous	Parallel-biased V	Mia hit Winona with a pea and Ian kicked her under the table.
	Result-biased V	Mia hit Winona with a pea and Ian scolded her.
	Result (*so*)	Mia hit Winona with a pea so Ian kicked her under the table.
Disambig (S)	Parallel-biased V	Mia hit Winona with a pea and Ian kicked her under the table. Mia didn't kick him in return, she simply...
Disambig (O)	Parallel-biased V	Mia hit Winona with a pea and Ian kicked her under the table. Winona didn't bother the other kids in return, she simply...
Unambig (S/O)	Parallel-biased V	Mia hit Harrison with a pea and Ian kicked (her/him) under the table.
Is it more likely that Ian kicked Mia or Winona? (Mia, Winona)		*How likely is the other option?* (1–4)

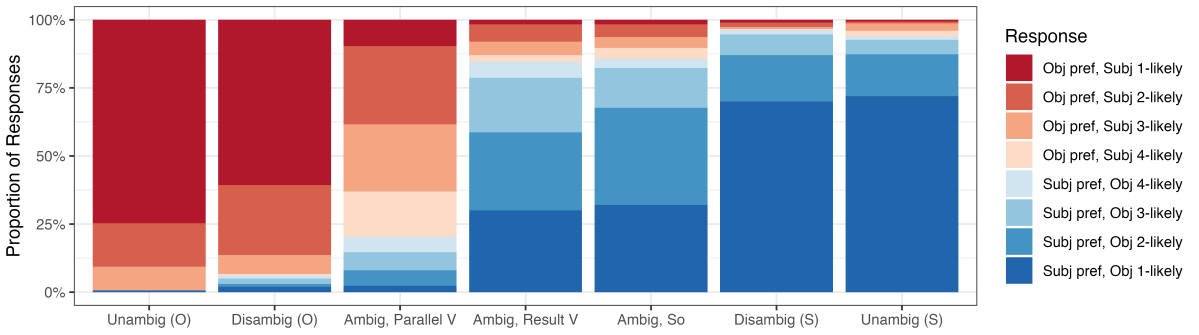

Figure 1: Results from Experiment 1.

coreference, like (1). Such discourses are uniquely interesting because of the complexity of the reanalysis they might engender, and because of the strength of their preferred interpretations. We present two experiments, first validating strong preferences for our set of narratives in a judgment study, and then examining their processing in self-paced reading. We find that when late information contradicts the preferred interpretation, readers experience a slowdown consistent with costly reanalysis. We take this to diagnose discourse reanalysis as a real phenomenon of incremental comprehension, with important consequence for formal theories.

2 Experiment 1: Interpretation preferences

In an initial experiment,[6] we collected preferred interpretations of potentially-ambiguous pronouns in 60 narratives like (1), with and without disambiguating continuations. Each narrative featured coordinated clauses with parallel syntactic structure $[S_1, S_2]$, where S_2 contained the critical pronoun in object position. The interpretation of such pronouns is biased towards an object antecedent due to parallelism (Smyth 1994), but this can be reversed, either by adding a preceding causal connective (Wolf, Gibson, and Desmet 2004), e.g. *so*, or by using a second verb biased more towards the causal relation Result than Parallel (Kehler et al. 2008). We elicited preferred antecedents for the object pronoun as a binary choice, followed by a 4-point Likert scale asking about the likelihood of the participant's dispreferred antecedent. These measures were compared across six key conditions (Table 1) to verify the contributions of the verb and connective, and further examine how preferred interpretations changed when we added an S_3 where the presupposition of *in return* disambiguated the previous pronoun.

[6]The materials, methods, and analytical choices for both experiments were pre-registered (Duff and Altshuler 2024a). Data and analysis scripts are available in our supplementary materials (Duff and Altshuler 2024b).

Table 2: Stimuli from Experiment 2. "/" indicates a chunk boundary. Critical regions in bold.

The kids/were misbehaving/at a fancy dinner.	
Ambig.	Mia hit Winona/with a pea/and/**Ian kicked her**/under the table.
Unambig.	Mia hit Harrison/with a pea/and/**Ian kicked her**/under the table.
so + Unambig.	Mia hit Harrison/with a pea/so/**Ian kicked her**/under the table.
Name	Winona hit Harrison/with a pea/and/**Ian kicked Mia**/under the table.
Mia didn't kick him/**in return**,/**she simply**/went to get more peas.	

We tested 30 native English speakers recruited on Prolific (Figure 1). Responses were combined to create an index of the direction and strength of interpretation preferences,[7] which was then analyzed using Bayesian ordinal mixed-effects regression.[8] Participants exhibited a strong preference (80%) for an object antecedent in the ambiguous, Parallel condition. This object antecedent preference was sharply reversed by causal-biased verbs, $\delta_{.95} = (-1.92, -1.24)$, by the causal connective *so*, $\delta_{.95} = (-2.05, -1.31)$, and, most strongly, by an S_3 disambiguating to a subject antecedent, $\delta_{.95} = (-3.34, -2.31)$.

These results replicate a strong object-antecedent bias in such narratives that is conditional on certain discourse coherence, and they provide a proof of concept that initially-preferred interpretations of ambiguous pronouns can be reversed by a subsequent sentence.

3 Experiment 2: Phrase-by-phrase reading times

From Experiment 1, we identified 32 narratives with the strongest parallelism bias in S_2 and the largest reversal given S_3, to investigate in a chunked, non-cumulative self-paced-reading experiment testing for costly reanalysis. We compared these 32 across four conditions (Table 2). In Ambiguous conditions, the critical object pronoun is expected to prefer an initial object antecedent due to establishing Parallel, but it is followed by an S_3 which requires a subject antecedent. In all other conditions, S_3 doesn't motivate any change in interpretation, because S_2 already makes a subject antecedent unambiguous, or lacks a pronoun. If comprehenders indeed commit to a firm interpretation of S_2 that they must then reanalyze, we should see particularly slow reading of S_3 in Ambiguous conditions. Materials were presented with 60 filler items of similar length, including 28 narratives where the discourse remained globally consistent with a Parallel relation and an object antecedent for the pronoun.

Reading times (Figure 2) from 96 participants (recruited as in Experiment 1) were analyzed using Bayesian log-normal mixed-effects regressions. We find that the final sentences were generally read slightly slower in the Ambiguous condition (red) than the other conditions, with our models estimating credible marginal slowdowns of about 15ms at *in return*, $\delta_{.95} = (1, 29)$, and *she simply*, $\delta_{.95} = (3, 25)$.[9] This result suggests that readers experienced some degree of

[7] An 8-point ordinal variable ranging from object-antecedent preferences where a subject antecedent was deemed unlikely, to subject-antecedent preferences where an object antecedent was deemed unlikely.

[8] All analyses for both experiments were conducted using `brms` (Bürkner 2017; Bürkner 2018) in R, using regularizing priors. In the text, we report 95% highest-density intervals over the posterior marginal differences between conditions extracted from our models, $\delta_{.95}$. When this interval excludes 0, we take the associated effect to be noteworthy. Note that δ in Experiment 1 is expressed in abstract model-internal units, not the units of the response scale; in Experiment 2, δ is the difference in reading times in milliseconds.

[9] Evidence for costly revision is strongest for the first trials of the experiment, and disappears by the last few, resembling adaptation effects well-known in work on syntactic garden paths (Prasad and Linzen 2021). We conjecture that participants may have learned to postpone discourse interpretation after experiencing that early interpretations often required reanalysis. An alternative possibility is that participants learned to anticipate the final subject-antecedent interpretation correctly, although we note that there is no evidence of a corresponding increase in discourse reanalysis for the fillers featuring late disambiguation to an object-antecedent interpretation.

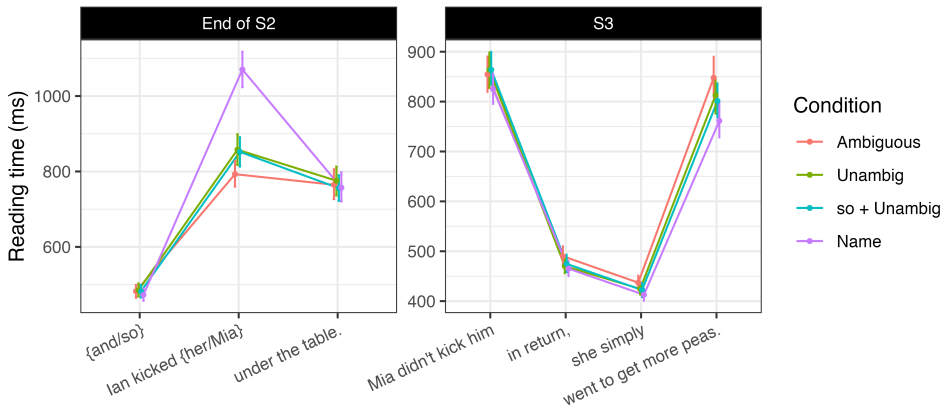

Figure 2: Results from Experiment 2.

costly reanalysis, revising an initial interpretation that yielded an object antecedent for *her*. This cost may include revision of the coherence relation between S_2 and S_3, from Parallel to Result; the cost of any such process cannot be untangled from pronominal revision.

Reading times on the core content of S_2 offer further insight. Besides costs for a new discourse referent (purple), we also see here an inversion of the later pattern; conditions with unambiguous subject antecedents (green, blue) were slower than Ambiguous (red), although this fell just below our standards for a notable contrast $\delta_{.95}$ = (-2, 57). This helps us rule out an alternative hypothesis for the later slow-down; ambiguous narratives were not globally slower as a function of some ambiguity-specific processing (cf. Stewart, Holler, and Kidd 2007; Creemers and Meyer 2022; see also Logačev and Vasishth 2016). This also cannot be a simple cost of processing an unambiguous pronoun (cf. Swets et al. 2008; Grant, Sloggett, and Dillon 2020), as unambiguous object-antecedent conditions in our filler items are not similarly slow. We take this instead as weak evidence that participants begin reading this region already expecting Parallel and thus parallelism in coreference, and comprehension is facilitated when the input matches this.[10]

4 Discussion

The reading data from Experiment 2 provides evidence for a small slowdown just when comprehenders encounter input that encourages a previously despised set of coreference and coherence choices. We take this to diagnose costly reanalysis of discourse interpretation, consistent with the idea that, at least in some cases, comprehenders make and revise firm pragmatic representations during incremental discourse comprehension, and do not somehow underspecify indefinitely.

Zooming out, the consequences of our results for theories of semantic representation depend on one's theoretical taste. We follow Brasoveanu and Dotlačil (2015) in considering two paths for how formal semantics may deal with the nature of incremental comprehension. On the first path, we could define a logic capable of representing and updating meaning incrementally. Among the theories available along this path, our results arbitrate for a dynamic semantics which permits selective non-monotonic update of coreference and other aspects of discourse meaning.

On the second path, we could maintain a simpler theory of representation, and shift the explanatory burden of incrementality onto our theory of processing. Along this path, we think our results help discriminate between two options for how a processing theory could treat discourse reanalysis. On the most extreme approach, we might imagine that discourse reanalysis requires the wholesale reconstruction of an alternative analysis permitted by the grammar, as carried out by a separate meta-system aware of the constraints from later input (a 'semantic

[10] As for implicit Results (Keenan, Baillet, and Brown 1984) and Explanations (Mak and Sanders 2013).

parser'). This would require no flexibility for discourse-pragmatic meaning in the semantic representations themselves; for instance, the conflation of pronouns and their antecedents that Haug (2014) objects to in DRT would not be a problem if the semantic parser re-parsed the entire relevant sequence with hindsight. On a more balanced approach, one could treat discourse reanalysis as an operation of a semantic parser, but using representations which do distinguish coreference (etc.) in a way which allows them to be selectively edited.[11] We take the specific reanalysis costs we observe in Experiment 2 to be far more compatible with the latter option. If our (surprisingly small) 15ms estimate is generalizable, the reanalysis of coreference seems to be rather less difficult than syntactic garden paths observed in self-paced reading, e.g., 50ms in Ferreira and Henderson (1990). This would be easier to understand if discourse reanalysis is unique in that it does not require complete reconstruction of the input—and even in a system where incrementality is largely handled outside the grammar, this would require a representation of meaning which permits selective revisions to discourse-pragmatic content.

Another open question across any of these approaches is *how often* comprehenders make and revise firm pragmatic representations; that is, we have established a case where we observe incremental interpretation and reanalysis, but is incremental interpretation the rule? Given the literature on processing coherence and coreference, we think we can best explain the diversity of previous findings if incrementality is flexible, and temporary underspecification may persist without preference in some cases. After all, other reading experiments have failed to observe costly discourse reanalysis, in cases of ambiguous coreference in a short narrative without connections to coherence (Stewart, Holler, and Kidd 2007), and in cases of ambiguous coherence without connections to coreference (Mak and Sanders 2013; Duff 2023), or with consequences for temporal ordering (Dickey 2001; Sasaki 2021). In finding some cost here, we join a few other experiments, where costly reanalysis was observed for coreference when an initial interpretation was more strongly cued by prominence as a repeated subject (Gordon and Scearce 1995), consistency with an explicit coherence relation (Mak and Sanders 2010), or consistency with world knowledge (Jones and Bergen 2024).[12] We think this split is likely the result of adaptively-timed interpretation, as would be permitted in logics like Asher and Lascarides (2003), or processing models like Karimi and Ferreira (2016) or Duff (2023), which lays out a larger argument that comprehenders learn to postpone interpretation for aspects of meaning which are contextually uncertain or practically irrelevant.

If costly reanalysis is a general matter of discourse comprehension, we predict similar effects when a strong coherence bias drives resolution of e.g. temporal ordering. To this end, our next experiments will test cases like (2), where S_1 and S_2 first appear to relate via Explanation, such that the event described by S_2 would be understood to precede the event described by S_1. S_3 would then force reanalysis to Result, ensuring the opposite event ordering.

(2) [Bob sued Liz.]$_{S_1}$ [She defaced his shop,]$_{S_2}$ [because his lawsuit was cruel.]$_{S_3}$

This reanalysis would be from a subordinating coherence relation to a coordinating one, in the sense of Asher and Vieu (2005), whereas our present study featured only coordinating relations. If coherence revision is a source of difficulty, as we hypothesize, this discourse structural shift may drive especially large reanalysis costs.

Acknowledgements. Thanks to Kelsey Sasaki, Christina Kim, Julian J. Schlöder, Adrian Brasoveanu and three anonymous reviewers for insightful input. The experiments were financially supported by the John Fell Fund, awarded to the authors by Oxford University Press.

[11] For a starting point, see Brasoveanu and Dotlačil's (2020) incremental DRT parser in ACT-R.

[12] Experiments with other designs are likewise split, compare Gernsbacher (1989) and Greene, McKoon, and Ratcliff (1992) vs. Arnold et al. (2000), Pyykkönen and Järvikivi (2010), and Rohde, Levy, and Kehler (2011).

References

Albrecht, Jason E. and Edward J. O'Brien (1993). "Updating a mental model: Maintaining both local and global coherence". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19, pp. 1061–1070. DOI: 10.1037/0278-7393.19.5.1061.

Allbritton, David W., Gail McKoon, and Roger Ratcliff (1996). "Reliability of prosodic cues for resolving syntactic ambiguity". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22.3, pp. 714–735. DOI: 10.1037/0278-7393.22.3.714.

Altshuler, Daniel and Dag T. T. Haug (2017). "The Semantics of provisional, temporal anaphors and cataphors". Talk presented at Deutsche Gesellschaft für Sprachwissenschaft, Saarland University, Saarbrücken.

— (forthcoming). *Literature as a formal language: The semantics and pragmatics of narrative garden path in 'Sylvie'*. Routledge.

Altshuler, Daniel and Christina S. Kim (2024). "Towards a typology of narrative frustration". In: *Topoi* 43, pp. 1193–1210. DOI: 10.1007/s11245-023-09942-1.

Arnold, Jennifer E. et al. (2000). "The rapid use of gender information: Evidence of the time course of pronoun resolution from eyetracking". In: *Cognition* 76, B13–B26. DOI: 10.1016/s0010-0277(00)00073-1.

Asher, Nicholas and Alex Lascarides (2003). *Logics of conversation*. Cambridge.

Asher, Nicholas and Laure Vieu (2005). "Subordinating and coordinating discourse relations". In: *Lingua* 115, pp. 591–610. DOI: 10.1016/j.lingua.2003.09.017.

Brasoveanu, Adrian and Jakub Dotlačil (2015). "Incremental and predictive interpretation: Experimental evidence and possible accounts". In: *Proceedings of SALT 25*, pp. 57–81.

— (2020). *Computational cognitive modeling and linguistic theory*. Springer. DOI: 10.1007/978-3-030-31846-8.

Bürkner, Paul-Christian (2017). "brms: An R package for Bayesian multilevel models using Stan". In: *Journal of Statistical Software* 80.1, pp. 1–28. DOI: 10.18637/jss.v080.i01.

— (2018). "Advanced Bayesian multilevel modeling with the R package brms". In: *The R Journal* 10, pp. 395–411. DOI: 10.32614/rj-2018-017.

Carlson, Katy et al. (2009). "Information structure expectations in sentence comprehension". In: *Quarterly Journal of Experimental Psychology* 62.1, pp. 114–139. DOI: 10.1080/17470210701880171.

Creemers, Ava and Antje S. Meyer (2022). "The processing of ambiguous pronominal reference is sensitive to depth of processing". In: *Glossa Psycholinguistics* 1.1, p. 3. DOI: 10.5070/G601166.

Dickey, Michael Walsh (2001). *The processing of tense*. Kluwer. DOI: 10.1007/978-94-010-0568-5.

Duff, John (2023). *On the timing of decisions about meaning during incremental comprehension*. Dissertation, UC Santa Cruz.

Duff, John and Daniel Altshuler (2024a). *Probing Parallel-to-Result reanalysis*. Preregistration hosted by the OSF, filed June 24, 2024. URL: https://osf.io/4sz79?view_only=bc75f487b79e47da9fe5c7f3765805e3.

— (2024b). *Reanalysis in discourse comprehension: Supplementary materials*. Repository hosted by the OSF. URL: https://osf.io/mk38v/?view_only=e01b727827b1465c95848f467b0713b1.

Duffy, Susan A., Robin K. Morris, and Keith Rayner (1988). "Lexical ambiguity and fixation times in reading". In: *Journal of Memory and Language* 27, pp. 429–446. DOI: 10.1016/0749-596X(88)90066-6.

Egg, Markus (2010). "Semantic underspecification". In: *Language and Linguistics Compass* 4.3, pp. 166–181. DOI: 10.1111/j.1749-818X.2010.00188.x.

Ferreira, Fernanda and John M. Henderson (1990). "Use of verb expectation in syntactic parsing: Evidence from eye movements and word-by-word self-paced reading". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 16.4, pp. 555–568. DOI: 10.1037/0278-7393.16.4.555.

Fodor, Janet Dean and Fernanda Ferreira (1998). *Reanalysis in sentence processing*. Springer. DOI: 10.1007/978-94-015-9070-9.

Frazier, Lyn and Keith Rayner (1982). "Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences". In: *Cognitive Psychology* 14.2, pp. 178–210. DOI: 10.1016/0010-0285(82)90008-1.

Gernsbacher, Morton A. (1989). "Mechanisms that improve referential access". In: *Cognition* 32, pp. 99–156. DOI: 10.1016/0010-0277(89)90001-2.

Gordon, Peter C. and Kimberly A. Scearce (1995). "Pronominalization and discourse coherence, discourse structure and pronoun interpretation". In: *Memory & Cognition* 23.3, pp. 313–323. DOI: 10.3758/BF03197233.

Grant, Margaret, Shayne Sloggett, and Brian Dillon (2020). "Processing ambiguities in attachment and pronominal reference". In: *Glossa* 5.1, p. 77. DOI: 10.5334/gjgl.852.

Greene, Steven B., Gail McKoon, and Roger Ratcliff (1992). "Pronoun resolution and discourse models". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 18.2, pp. 266–283. DOI: 10.1037/0278-7393.18.2.266.

Haug, Dag T. T. (2014). "Partial dynamic semantics for anaphora: Compositionality without syntactic coindexation". In: *Journal of Semantics* 31, pp. 457–511. DOI: 10.1093/jos/fft008.

Hobbs, Jerry R. (1979). "Coherence and coreference". In: *Cognitive Science* 3, pp. 67–90. DOI: 10.1207/s15516709cog0301_4.

Jaspars, Jan and Megumi Kameyama (1998). "Discourse preferences in dynamic logic". In: *Computing natural language*. Ed. by Atocha Aliseda, Rob van Glabbeek, and Dag Westertahl. CSLI, pp. 67–96.

Jones, Cameron R. and Benjamin Bergen (2024). "Does word knowledge account for the effect of world knowledge on pronoun interpretation?" In: *Language and Cognition*. Advance online publication. DOI: 10.1017/langcog.2024.2.

Kameyama, Megumi (1996). "Indefeasible semantics and defeasible pragmatics". In: *Quantifiers, deduction, and context*. Ed. by Makoto Kanazawa, Christopher Piñón, and Henriette de Swart. CSLI, pp. 111–138.

Kamp, Hans and Uwe Reyle (1993). *From discourse to logic: Introduction to model theoretic semantics of natural language, formal logic and Discourse Representation Theory*. Kluwer. DOI: 10.1007/978-94-017-1616-1.

Karimi, Hossein and Fernanda Ferreira (2016). "Good-enough linguistic representations and online cognitive equilibrium in language processing". In: *Quarterly Journal of Experimental Psychology* 69.5, pp. 1013–1040. DOI: 10.1080/17470218.2015.1053951.

Keenan, Janice M., Susan D. Baillet, and Polly Brown (1984). "The effects of causal cohesion on comprehension and memory". In: *Journal of Verbal Learning and Verbal Behavior* 23, pp. 115–126. DOI: 10.1016/S0022-5371(84)90082-3.

Kehler, Andrew et al. (2008). "Coherence and coreference revisited". In: *Journal of Semantics* 25.1, pp. 1–44. ISSN: 0167-5133. DOI: 10.1093/jos/ffm018.

Lascarides, Alex and Nicholas Asher (1993). "Temporal interpretation, discourse relations and commonsense entailment". In: *Linguistics and Philosophy* 16, pp. 437–493. DOI: 10.1007/BF00986208.

Logačev, Pavel and Shravan Vasishth (2016). "A multiple-channel model of task-dependent ambiguity resolution in sentence comprehension". In: *Cognitive Science* 40, pp. 266–298. DOI: 10.1111/cogs.12228.

Mak, Willem M. and Ted J. M. Sanders (2010). "Incremental discourse processing: How coherence relations influence the resolution of pronouns". In: *The linguistics enterprise: From knowledge of language to knowledge in linguistics*. Ed. by Martin Everaert et al. John Benjamins, pp. 167–182.

— (2013). "The role of causality in discourse processing: Effects of expectation and coherence relations". In: *Language and Cognitive Processes* 28.9, pp. 1414–1437. ISSN: 0169-0965. DOI: 10.1080/01690965.2012.708423.

Prasad, Grusha and Tal Linzen (2021). "Rapid syntactic adaptation in self-paced reading: Detectable, but only with many participants". In: *Journal of Experimental Psychology: Learning, Memory, and Cognition* 47.7, pp. 1156–1172. DOI: 10.1037/xlm0001046.

Pyykkönen, Pirita and Juhani Järvikivi (2010). "Activation and Persistence of Implicit Causality Information in Spoken Language Comprehension". In: *Experimental Psychology* 57.1, pp. 5–16. DOI: 10.1027/1618-3169/a000002.

Rohde, Hannah, Roger Levy, and Andrew Kehler (2011). "Anticipating explanations in relative clause processing". In: *Cognition* 118.3, pp. 339–358. DOI: 10.1016/j.cognition.2010.10.016.

Sasaki, Kelsey (2021). *Components of coherence*. Dissertation, UC Santa Cruz.

Smyth, Ron (1994). "Grammatical determinants of ambiguous pronoun resolution". In: *Journal of Psycholinguistic Research* 23.3, pp. 197–229. DOI: 10.1007/BF02139085.

Stewart, Andrew J., Judith Holler, and Evan Kidd (2007). "Shallow processing of ambiguous pronouns: Evidence for delay". In: *Quarterly Journal of Experimental Psychology* 60.12, pp. 1680–1696. DOI: 10.1080/17470210601160807.

Swets, Benjamin et al. (2008). "Underspecification of syntactic ambiguity: Evidence from self-paced reading". In: *Memory & Cognition* 36.1, pp. 201–216. DOI: 10.3758/mc.36.1.201.

Syrett, Kristen, Georgia Simon, and Kirsten Nisula (2012). "Prosodic disambiguation of scopally ambiguous sentences". In: *Proceedings of NELS 43*, pp. 141–152.

Wolf, Florian, Edward Gibson, and Timothy Desmet (2004). "Discourse coherence and pronoun resolution". In: *Language and Cognitive Processes* 19.6, pp. 665–675. DOI: 10.1080/01690960444000034.