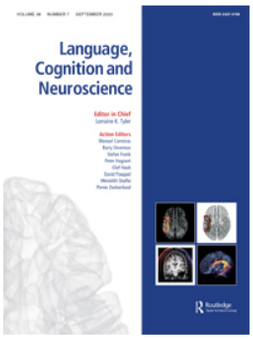

Language, Cognition and Neuroscience

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/plcp21

Effects of referential structure on pronoun interpretation

Jina Song & Elsi Kaiser

To cite this article: Jina Song & Elsi Kaiser (2023): Effects of referential structure on pronoun interpretation, Language, Cognition and Neuroscience, DOI: 10.1080/23273798.2023.2250481

To link to this article: https://doi.org/10.1080/23273798.2023.2250481

LANGUAGE, COGNITION AND NEUROSCIENCE
https://doi.org/10.1080/23273798.2023.2250481

Routledge
Taylor & Francis Group

REGULAR ARTICLE

Effects of referential structure on pronoun interpretation

Jina Song and Elsi Kaiser

Department of Linguistics, University of Southern California, Los Angeles, CA, USA

ABSTRACT

Pronoun interpretation is guided by various factors. While most previously-investigated factors involve properties occurring before the pronoun, less attention has been paid to properties of the pronoun-containing clause. We investigate whether pronoun interpretation is influenced by the referential structure of the pronoun-containing clause (i.e. whether another referent from the preceding clause is mentioned), which contributes to discourse coherence. We report three experiments showing referential structure effects: whether subject-position pronouns are ultimately interpreted as referring to the preceding subject or object depends on whether the clause contains another pronoun (e.g. *she called Lisa* vs. *she called her*). More specifically, subject-position pronouns exhibit a stronger object preference when only one of the prior antecedents is mentioned, compared to when both are mentioned. We show that this effect is separate from effects of verb semantics and cannot be reduced to semantic or syntactic parallelism effects. Implications for models of pronoun resolution are discussed.

ARTICLE HISTORY
Received 6 July 2022
Accepted 12 August 2023

KEYWORDS
Pronoun interpretation; discourse coherence; implicit causality; transitivity; reference resolution

Introduction

Pronoun interpretation is one of the core questions for theories of communication and has drawn the attention of many researchers. As third-person pronouns (e.g. *she, he*) provide very little semantic information about their referents, their interpretation must rely on other sources of information. Despite being referentially underspecified, pronouns are known to be interpreted efficiently and rapidly. Previous psycholinguistic work has investigated reference resolution from many perspectives. A widely-shared assumption is that reduced anaphoric expressions like pronouns tend to refer to highly salient entities, i.e. entities that are prominent in comprehenders' mental models of discourse (e.g. Ariel, 1990, 1994; Givón, 1983; Gundel et al., 1993). Arnold et al. (2013, p. 404) define salient entities as "those that are more activated or accessible in people's minds at that point in the discourse". Various factors, including syntactic role, order of mention, discourse status and thematic roles, have been argued to influence the salience of referents, and therefore to influence the likelihood of a particular referent being interpreted as the antecedent of a subsequent pronoun. For example, prior work shows that – other things being equal – third-person pronouns in subject position tend to be interpreted as referring to preceding subjects (e.g. Bosch & Umbach, 2007; Crawley et al., 1990; Crawley & Stevenson, 1990; Grober et al., 1978; Kaiser & Trueswell, 2008), linearly first-mentioned entities (e.g. Ellert, 2013; Gernsbacher et al., 1989; Gernsbacher & Hargreaves, 1988), and more generally to salient entities (e.g. Ariel, 1990, 1994; Arnold, 1998; Colonna et al., 2012, 2014; Givón, 1983; Kaiser, 2011a). Prior work also shows effects of antecedents' thematic roles. Most relevantly for us, in causal contexts with implicit causality (IC) verbs, subsequent pronouns tend to be interpreted as referring to Stimulus referents more often than Experiencers (e.g. Au, 1986; Bott & Solstad, 2014; Caramazza et al., 1977; Ferstl et al., 2011; Hartshorne & Snedeker, 2013; McKoon et al., 1993; Patterson et al., 2022; Rudolph & Försterling, 1997).

The factors mentioned so far involve properties that occur before the pronoun itself is encountered. Factors such as grammatical role, thematic role and linear position are all characteristics of potential antecedents, and thus constitute information that occurs in a clause that precedes the pronoun-containing clause. Most psycholinguistic experiments on pronoun resolution have tended to focus on how pronoun interpretation is guided by this kind of preceding information. Given the highly incremental nature of language processing, this is not unexpected. The flip side is that potential effects of information occurring later in the pronoun-containing clause have not been systematically investigated in psycholinguistic research to the same extent. Ultimately, a satisfactory model of pronoun

CONTACT Jina Song jinas@usc.edu Department of Linguistics, University of Southern California, 3601 Watt Way, Los Angeles, CA 90089-1693, USA
Supplemental data for this article can be accessed online at https://doi.org/10.1080/23273798.2023.2250481.

interpretation should capture all relevant factors. Thus, if our aim is to account for the mechanisms that explain how comprehenders ultimately interpret pronouns, it is necessary to gain a better understanding of how different kinds of information in the pronoun-containing clause guide pronoun interpretation.

In classic work, Winograd (1972) pointed out that subsequent semantic cues in the pronoun-containing clause can guide interpretation of subject-position pronouns (e.g. *The city council denied the demonstrators the permit because they {feared/advocated} violence*). The contribution of information in the pronoun-containing clause is also acknowledged by coherence-based approaches (e.g. Hobbs, 1979; Kaiser, 2011b; Kehler, 2002; Kehler & Rohde, 2013; Song & Kaiser, 2020a; Wolf et al., 2004). In the studies reported in this paper, we aim to further our understanding of how information in the pronoun-containing clause guides pronoun interpretation. Specifically, we systematically investigate whether and how a foundational referential property of the pronoun-containing clause – namely, who else is or isn't mentioned in that clause – influences pronoun interpretation. We use an offline task and focus on end-of-sentence interpretations, because our aim is to establish whether people's final pronoun interpretation is influenced by referential structure. We regard this work as providing a necessary foundation for future investigations using real-time methods.

Referential properties of the pronoun-containing clause

Our work focuses on the interpretational consequences of a specific type of information in the pronoun-containing clause, namely its *referential structure*. For example, in (1), the interpretation of the subject-position pronoun *she* is influenced by whether no other referents are mentioned in the rest of the clause (1a), whether a new referent is introduced (1b) or whether another third-person pronoun occurs in the same clause (1c).

(1)	a. Lisa surprised Mary because **she** frobbed.	(One Pronoun)
	b. Lisa surprised Mary because **she** frobbed Kate.	(Pronoun + Name)
	c. Lisa surprised Mary because **she** frobbed her.	(Two Pronouns)

Before considering our specific predictions, we review work on Centering Theory (e.g. Grosz et al., 1995; Walker et al., 1998), which provides the starting point for why one might expect effects of referential structure on pronoun resolution. (Although the *Referential Structure Hypothesis* that we propose is inspired by Centering Theory, it is not an application or test of Centering and diverges from it in various ways; for example, we predict effects of referential structure even in two-clause sequences that lack a pre-established topic. Our claims and conclusions are not reliant on the validity of Centering Theory.)

What is relevant for us is the general idea that there is a preference to interpret anaphoric expressions so that the transitions between different utterances (typically finite clauses) are as *coherent* as possible (e.g. Grosz et al., 1995; Walker et al., 1998). Simplifying somewhat, the core idea is as follows: the less the most salient entity changes from one clause to the next, the more coherent the transition is. In other words, keeping one's attention "centered" on the same entity across clauses yields a maximally coherent discourse, and thus shapes comprehenders' pronoun interpretation biases. Like many other approaches, Centering Theory posits that English salience is largely defined in terms of grammatical role, with referents in subject position being more salient than those in object position (e.g. Grosz et al., 1995; Walker et al., 1998).

Our work builds on the Centering-based idea that a sequence of two clauses (let's call them Clause 1 and Clause 2) is highly coherent when, out of all Clause 1 referents re-mentioned in Clause 2, the most salient member of this set is mentioned in the most salient position in Clause 2 (namely, as Clause 2's subject). Conversely, mentioning an entity that was salient in Clause 1 in a downgraded, lower-salience position in Clause 2 – for example, demoting the subject of Clause 1 to the object position of Clause 2 – yields a less coherent transition. Thus, a two-pronoun clause (e.g. *she verbed her*) is expected to have a strong preference for *she* to refer to the preceding Clause 1 subject and *her* to refer to the preceding Clause 1 object. This is because the other interpretation, where *she* refers to the preceding object and *her* to the preceding subject has two undesirable properties: (i) it fails to realise the most salient member of the re-mentioned Clause 1 referents in the most salient position of Clause 2 and also (ii) demotes this previously-salient referent to a lower-salience position.

This generates the prediction that, in a Two-Pronoun configuration, the subject-pronoun *she* is expected to prefer the preceding subject over the preceding object as its antecedent, because the former yields a more coherent transition than the latter. This prediction obtains even when the two clauses are not linked by a parallelism/similarity relation.

In contrast, in a clause with only one pronoun, where only one of the Clause 1 referents is re-mentioned (e.g. *she verbed* or *she verbed Kate*), we expect that the subject-position pronoun *she* can refer to either the preceding subject or the preceding object, without any difference in the coherence of the resulting transition.

This is because now, only one of the entities of Clause 1 is re-mentioned in Clause 2, and as long as that referent is realised as the subject of Clause 2, the transition is highly coherent. In this regard, whether both or only one of the Clause 1 referents are re-mentioned in Clause 2 plays a key role.

In essence, according to this line of thinking, the *referential structure* of Clause 2 – in other words, whether one or both of the Clause 1 referents are re-mentioned in Clause 2 – plays a key role in guiding pronoun interpretation. Clauses that only mention one of the referents from the preceding clause (the One-Pronoun and Pronoun + Name versions in (1a,b)) are expected to elicit *relatively higher rates of object interpretations* than clauses that mention both referents from the preceding clause (the Two-Pronoun version in (1c)). We use the term ***Referential Structure Hypothesis*** for this prediction, and report a series of three experiments testing it, using different verb classes to assess its empirical robustness and generalizability.

It is important to note that unlike parallelism-based accounts (e.g. Chambers & Smyth, 1998; Smyth, 1994; Stevenson et al., 1995), the Referential Structure Hypothesis does not make any reference to the semantic relation of parallelism (and does not require the two clauses to be linked via a parallel relation), and instead simply focuses on the whether one or both of the Clause 1 referents are re-mentioned in Clause 2. This approach is entirely compatible with additional effects of semantic coherence relations, but it also predicts that we should find differences between Two-Pronoun configurations on the one hand, and One-Pronoun configurations on the other hand, even when the semantic/pragmatic relation between the clauses is an explanation relation as in (1), and not a parallelism relation.

Furthermore, the predictions of the Referential Structure Hypothesis are distinct from those of a purely syntactic-parallelism based account. If each individual pronoun simply seeks an antecedent in a syntactically parallel position, then subject-position pronouns in the Pronoun + Name and Two-Pronoun conditions (ex.1b,1c) should both seek a transitive subject as their antecedent. In other words, a syntactic-parallelism account does not directly predict that the subject pronoun in Pronoun + Name and Two-Pronoun conditions would be interpreted differently – contrary to the Referential Structure Hypothesis.

In sum, empirically speaking, the Referential Structure Hypothesis states that interpretation of the subject-position pronoun depends on whether an object-position pronoun is also present, independent of the semantic coherence relation at play. This dependency is not predicted by syntactic-parallelism accounts, at least not without additional assumptions about dependencies between subject and object pronoun interpretation that would essentially end up turning such accounts into versions of the Referential Structure Hypothesis. Theoretically speaking, the Referential Structure Hypothesis attributes these effects to a preference to maximise discourse coherence.

Psycholinguistic work on effects of referential structure

To the best of our knowledge, the existence of referential structure effects on pronoun resolution has received little attention in psycholinguistic experiments. Generally speaking, prior experimental work on pronoun interpretation has mostly tested pronoun use and interpretation after transitive sentences mentioning two animate referents (e.g. Bott & Solstad, 2014; Ferstl et al., 2011; Hartshorne & Snedeker, 2013; Hartshorne et al., 2015; Kehler et al., 2008; Patterson et al., 2022; Runner & Ibarra, 2016 ; Schumacher et al., 2017; see also Kehler et al., 2008; Rohde et al., 2006; Stevenson et al., 1994; Ueno & Kehler, 2016 on ditransitive sentences). However, these studies and others usually focus only on the interpretation and use of subject pronouns (e.g. Arnold, 1998; Colonna et al., 2012, 2014; Cowles, 2007 ; Hartshorne & Snedeker, 2013; Kaiser, 2011a; Koornneef & Sanders, 2013; Rohde et al., 2006) – in other words, they do not systematically test whether and how the presence of a subsequent pronoun after a subject-position pronoun would influence pronoun resolution.

An exception comes from work on parallelism effects, which has tested sentences with both a subject-position pronoun and a second non-subject-position pronoun (e.g. Chambers & Smyth, 1998; Gordon et al., 1993). It is important to note that these studies have largely focused on cases of semantic parallelism, i.e. contexts where the two clauses are not only syntactically but also semantically parallel to each other. As we discuss below, the Referential Structure Hypothesis that we test in this paper is not restricted to semantically-parallel contexts.

However, there is prior work by Gordon et al. (1993) and Kaiser (2009) that relates to the Referential Structure Hypothesis more directly. Although Gordon et al. (1993) did not focus on the kinds of referential structure effects investigated in the present paper, looked mostly at unambiguous pronouns, and did not probe pronoun resolution (they focused on the repeated name penalty), their reading-time data provides evidence for many of the core tenets of Centering Theory. In other

related work, Kaiser (2009) used a production-based fragment-completion task and found that pronoun interpretation is modulated by whether the other referent from the preceding sentence is mentioned later in the pronoun-containing sentence. Kaiser suggests that if a less salient referent from the preceding sentence is *promoted* by being mentioned in the subsequent subject-position with a pronoun, then participants will avoid mentioning a highly-salient referent from the preceding sentence later in the pronoun-containing sentence. However, this work only tested sentences with contrastive focus (i.e. a particular information-structural configuration), and did not systematically assess the impact of different kinds of referential structures or verb classes. Thus, although it provides valuable foundational data, more systematic investigation is needed to assess the validity of the Referential Structure Hypothesis.

Our work builds on Kaiser (2009) but goes beyond it (i) by directly testing how and whether pronoun interpretation varies depending on differences in referential structure (One-Pronoun, Two-Pronoun, and Pronoun + Name configurations), (ii) by investigating referential structure effects in contexts with implicit causality verbs, which allow us to assess whether well-known implicit causality effects replicate with our task while also building carefully on prior norming data that allows us to control verbs' referential biases, and (iii) by testing different verb classes in a systematic way (psych verbs and action verbs) to see whether referential structure effects generalise across verb classes with different thematic roles.

Effects of implicit causality

In our studies on the Referential Structure Hypothesis, we manipulated the implicit causality of the transitive verbs in the first clause of sentences like (1). It is well known that subject-position pronouns following certain verbs, in the presence of an explanation (*because*) relation, exhibit systematic interpretational biases depending on the verb (e.g. Bott & Solstad, 2014; Caramazza et al., 1977; Ferstl et al., 2011; Hartshorne & Snedeker, 2013; Hartshorne et al., 2015; Koornneef & van Berkum, 2006; McKoon et al., 1993; Rudolph & Försterling, 1997). Researchers have identified subject-biased verbs (IC1 verbs) that elicit subject interpretations (e.g. *X bothered Y because she … , X impressed Y because she …*) as well as object-biased verbs (IC2 verbs) that elicit object interpretations (e.g. *X worshipped Y because she … , X criticized Y because she …* .).

Even beyond explanation relations, Kehler's work (e.g. Kehler, 2002; Kehler et al., 2008) led to an increase in psycholinguistics studies recognising the importance of controlling for coherence relations (e.g. Kaiser, 2011b on Result vs. Narrative relations; Kertz et al., 2006; Kehler et al., 2008; Wolf et al., 2004 on Cause/Effect vs. Parallel relations; Rohde, 2008; Rohde & Kehler, 2008; Ueno & Kehler, 2016 on various coherence relations including Cause/Effect, Elaboration and Occasion).

The studies reported here focus on the explanation relation and included an implicit causality manipulation in our studies for two main reasons. First, this approach allows us to assess whether classic implicit causality effects replicate with the new task that we developed for the purpose of our experiments. If we find the same IC effects as previous studies, this would provide evidence that our method provides meaningful data about pronoun interpretation, and can thus be used to test the Referential Structure Hypothesis. Second, implicit causality contexts allow us to test explanation relations (indicated by the connective *because*) in a natural way – in other words, we can test pronoun interpretation in contexts that clearly do not involve semantic parallelism relations between clauses. This is important because we want to assess whether, as predicted by the Referential Structure Hypothesis, the predicted results occur independently of semantic parallelism effects (see Chambers & Smyth, 1998; Smyth, 1994; Stevenson et al., 1995 on semantic parallelism on pronoun interpretation). Demonstrating this would show that the Referential Structure Hypothesis cannot be subsumed under semantic parallelism effects and needs to be acknowledged as a distinct phenomenon.

In what follows, we report the outcomes of three experiments. Experiment 1 tests for effects of referential structure and implicit causality by comparing pronoun interpretation in intransitive and transitive continuations, containing one and two pronouns respectively, following clauses containing Stimulus-Experiencer verbs (SE, e.g. *astonished, surprised*) and Experiencer-Stimulus verbs (ES, e.g. *feared, believed*).

Experiment 2 addresses a potential confound in Experiment 1 by eliminating differences in transitivity and comparing pronoun interpretation in transitive clauses with one vs. two pronouns, again after clauses with Stimulus-Experiencer (SE) and Experiencer-Stimulus (ES) verbs. Experiment 3 broadens the domain of investigation beyond SE/ES verbs and tests whether referential structure effects also obtain in configurations with verbs that have agentive subjects, using Agent-Patient verbs (AP, e.g. *cheated, followed*) and Agent-Evocator verbs (AE, e.g. *criticized, welcomed*). As we will see, these three studies provide clear evidence in favour of

our claim that the referential structure of the entire pronoun-containing clause influences participants' final (offline) interpretation of pronouns in subject position.

Effects of thematic role

Our studies test four different kinds of verb types (IC1-biased SE verbs, IC2-biased ES verbs, IC1-biased AP verbs and IC2-biased AE verbs). Prior work on pronoun processing has also tested a variety of verb types, including AP verbs (e.g. Bott & Solstad, 2014; Fuchs & Schumacher, 2020; Patterson et al., 2022; Schumacher et al., 2016, 2017; Stevenson et al., 1994), transfer-of-possession verbs (e.g. Kehler & Rohde, 2013; Kehler et al., 2008; Rohde et al., 2006, on Japanese see Ueno & Kehler, 2016), SE/ES verbs (e.g. Bott & Solstad, 2014; Patterson et al., 2022; Stevenson et al., 1994), as well as Experiencer (dative)-Patient(nominative) verbs in German (e.g. Fuchs & Schumacher, 2020; Patterson et al., 2022; Schumacher et al., 2016, 2017). In our studies, we chose to focus specifically on implicit-causality configurations involving explanation relations (signalled by a *because* connective), while testing four different verb types that differ in their thematic-role properties. This allows us to test (i) whether implicit causality effects replicate in our study, (ii) whether referential structure effects obtain in a context that clearly does not involve semantic parallelism and (iii) whether referential structure effects persist in configurations involving different thematic roles. (The Evocator role is, semantically, equivalent to Patient, see e.g. Au, 1986; Ferstl et al., 2011; Rudolph & Försterling, 1997; see also footnote 4.)

Pronoun resolution is known to be sensitive to antecedents' thematic roles. According to a standard view of the Thematic Hierarchy, the roles relevant for the claims being made in this paper can be ranked as follows: *Agent* > *Experiencer* > *Stimulus* > *Patient* (e.g. Belletti & Rizzi, 1988; Bresnan & Kanerva, 1989; Foley, 2005; Foley & Van Valin, 1984; Giorgi, 1984; Grimshaw, 1990; Speas, 1990; Van Valin, 2001; Van Valin & LaPolla, 1997).

Existing work on reference resolution suggests that, other things being equal, pronouns typically tend to prefer antecedents that are higher-ranked on the thematic hierarchy over those that are lower-ranked, which fits with the idea that the ranking of thematic roles reflects entities' thematic prominence (see e.g. Rappaport Hovav & Levin, 2015 for related discussion).

As regards the four roles included in our studies (Agent vs. Patient and Experiencer vs. Stimulus), prior work largely agrees that the Agent role is thematically more prominent than the Patient role, and the Experiencer role is more thematically prominent than the Stimulus role (see also Rappaport Hovav and Levin (2015) for a review). This is often attributed to factors such as sentience, with thematic roles that are strongly associated with sentience (e.g. Agent, Experiencer) being more prominent than thematic roles that can be occupied by non-sentient referents (e.g. Patient, Stimulus).

What's relevant for us are the differences in the syntax-semantics mappings of different kinds of IC verbs. Experiments 1 and 2 test Stimulus-Experiencer (SE) IC1 and Experiencer-Stimulus (ES) IC2 verbs, which differ in that with ES verbs, the thematic and syntactic roles are *aligned* (the highest-ranked thematic role is in the highest-ranked syntactic position; the Experiencer is the subject). Thus, thematic and syntactic prominence are aligned. In contrast, with SE verbs thematic and syntactic prominence are *misaligned* (the highest-ranked thematic role, the Experiencer, is not in the highest-ranked syntactic position). (See also Fuchs & Schumacher, 2020; Schumacher et al., 2016, 2017, as well as Do & Kaiser, 2022 and Ferreira, 1994 for production-based work). Generally speaking, prior work suggests that in cases of misalignment, interpretation preferences are less stable (see e.g. Schumacher et al., 2017 on dative experiencers in German).

However, not all IC2 verbs exhibit this kind of misalignment. In Experiment 3, we test Agent-Patient (AP) IC1 verbs and Agent-Evocator (AE) IC2 verbs. Here, regardless of whether the IC bias is towards the subject or the object, thematic and syntactic prominence are *aligned* (the highest-ranked thematic role, the Agent, is also in the highest-ranked syntactic position). As we will see, these differences interact with referential structure effects.

Experiment 1

Method

Participants

Forty-five native English-speaking adults (aged 18 and over), recruited from Amazon MTurk, took part. We excluded five participants because they were not born in the U.S. or were not native speakers of English. (No participants were excluded for poor performance on catch trials, as everyone met the pre-specified threshold: 4 or more correct out of 6 catch trials.) Thus, forty participants (22 self-identified as female, 19 as male) were included in the final analyses.

Design and materials

Experiment 1 employed a 2 × 2 design manipulating (i) the type of IC verb in the first clause (IC1 vs. IC2) and (ii) the nature of the referential structure in the second clause (One-Pronoun vs. Two-Pronoun), as illustrated in

Table 1. An example item from Experiment 1.

Verb Type	Referential structure	Examples
IC1 (SE)	Two-Pronoun	Stacy discouraged Hanna because she kreeged her.
	One Pronoun	Stacy discouraged Hanna because she kreeged.
IC2 (ES)	Two-Pronoun	Stacy believed Hanna because she kreeged her.
	One Pronoun	Stacy believed Hanna because she kreeged.

Table 1. The two clauses were connected by *because*, signalling an explanation relation. The study had 24 target items.

In the first clause, verb semantics was manipulated by using Stimulus-Experiencer (SE) verbs (which have an IC1 bias) vs. Experiencer-Stimulus (ES) verbs (which have an IC2 bias). To ensure that the SE verbs were indeed subject-biased and that the ES verbs were object-biased in the expected way, we used the norms from Hartshorne and Snedeker (2013) and Ferstl et al. (2011) to select twenty-four SE verbs (mean subject bias = 67.4%, SD = 13.6) and twenty-four ES verbs (mean object bias = 76.3%, SD = 11.7).

The second clause manipulated referential structure and contained one or two pronouns. As shown in Table 1, the subject of the nonce verb (e.g. *kreeged*) is always a pronoun (*he* or *she*). In the One-Pronoun condition the verb is intransitive; in the Two-Pronoun condition the verb is transitive and followed by an object pronoun (*him* or *her*). All the verbs used in the second clause were nonce verbs (e.g. *daxed, kreeged, frobbed*).[1] This allowed us to minimise potential effects of these verbs' semantics on pronoun interpretation, a question we do not aim to test in the present work. Each nonce verb was used once. Sentences contained two-same gender names (two male or two female referents), such that the pronouns were ambiguous.

In addition to the 24 targets, the study included 36 fillers. The fillers were of various types, including ambiguous relative-clause structures and negative stripping structures (e.g. *X verbed Y but not Z*). In addition, six of these fillers were unambiguous catch trials that allowed us to check whether participants were paying attention to the task. The target and filler items were intermixed and presented using a Latin-Square design. The complete list of target items can be found in Appendix A.

Procedure

We used a picture-writing task, where participants saw a sentence coupled with a schematised picture depicting the underlined part of the sentence. An example is in Figure 1. In target items, the underlined part was the critical pronoun-containing clause. The participants' task was to type the names of the characters into the textboxes. This reveals how they interpret the pronoun (s). The study was conducted online using Qualtrics (Provo, UT, 2019).

We opted not to measure reaction times, because (i) we did not want to put participants under time pressure and inadvertently encourage shallow processing of pronouns (see e.g. Creemers & Meyer, 2022; Stewart et al., 2007) and because (ii) variation in typing speed is likely to yield highly variable reaction time data in this kind of complex task.

The target pictures included one or two stick figures (depicting the characters involved in the event), and a nonce verb for the action in the pronoun-containing clause (e.g. *kreeged*). For one-character pictures (Figure 1(b); intransitive One-Pronoun conditions), the character in the picture is the one who does the action (Agent). Thus, with two-character images, the nonce verbs are transitive and with one-character images they are intransitive.

For pictures with two characters (Figure 1(a); transitive Two-Pronoun conditions), the character at the beginning of the arrow does the action (Agent), and the character at the end of the arrow is the undergoer (Patient/Theme). This was explained to participants as part of the instructions. The left/right directionality of the arrows as well as the location of the agents and patients/themes (left vs. right) in the two-character pictures were counterbalanced.

The task was to type a name into each text box such that the picture fits with the underlined part of the sentence. For example, in Figure 1(a), if a participant interprets the second clause as *Stacy kreeged Hanna*, they should type *Stacy* in the left textbox and *Hanna* in the right textbox. Thus, participants' responses reveal how they interpret the pronouns.

Data processing

Responses were coded according to which of the referents in the first clause (subject or object) was chosen as the referent of the *subject-position* pronoun: We analyse how often this subject-position pronoun in the second clause is interpreted as referring to the *object* of the first clause by coding object choices as 1 and subject choices as 0. When participants typed in a name that referred to neither the subject nor object of the preceding clause, the trial was marked as N/A (0.7% of all responses) and excluded from analysis.

Because there are only two possible antecedents, looking at how the *subject*-position pronoun is

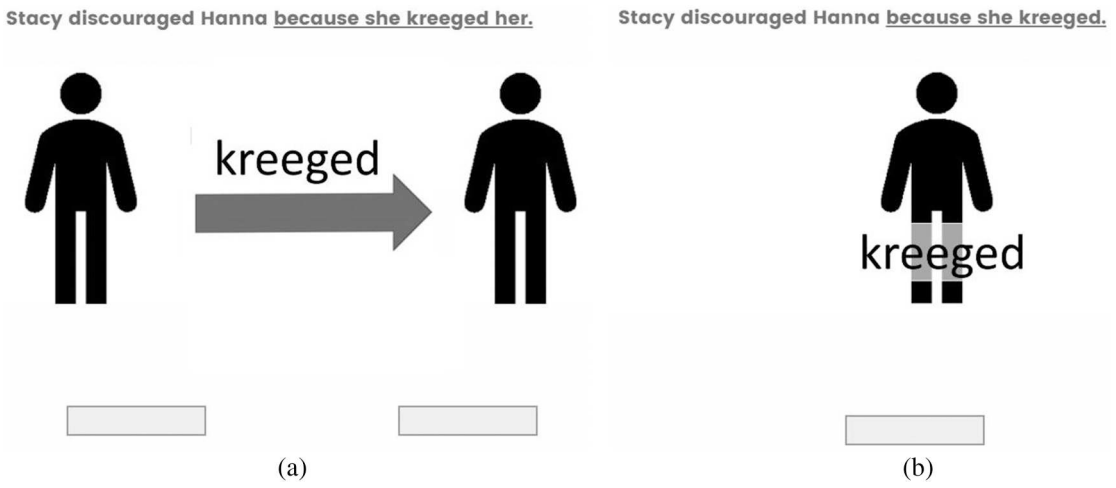

Figure 1. (a) Experiment 1: Sample Two-Pronoun item. (b) Experiment 1: Sample One-Pronoun item.

interpreted also provides information about how the object pronoun (if present) is interpreted: If the subject-position pronoun refers to the preceding subject, the object-position pronoun refers to the preceding object; if the subject-position pronoun refers to the preceding object, the object-position pronoun refers to the preceding subject.

Predictions

The main aim of our study is to test the Referential Structure Hypothesis, according to which pronoun resolution is affected by the referential properties of the pronoun-containing clause as described above. We test for referential structure effects in the presence of previously-observed implicit causality/verb semantic effects, to check whether our method reliably detects the expected verb bias effects. One possible outcome is that pronoun interpretation is guided only by the implicit causality properties of the preceding verb – i.e. no effects of referential structure. In this case, we expect that subject-position pronouns will be more likely to refer to the preceding object antecedent in the ES (IC2) condition than in the SE (IC1) condition, and that subject-position pronouns will pattern alike in One-Pronoun and Two-Pronoun conditions.

However, if pronoun interpretation is sensitive to the referential structure of the pronoun-containing clause, we expect differences between the Two-Pronoun and One-Pronoun conditions. The Referential Structure Hypothesis predicts that in One-Pronoun conditions, the subject-position pronoun is more likely to refer to a preceding object antecedent (and less likely to refer to a preceding subject antecedent) compared to Two-Pronoun conditions. This is because, in the Two-Pronoun conditions (when both entities from the preceding clause are mentioned), the most coherent transition results from the subject-position pronoun referring to the *subject* antecedent. In contrast, in the One-Pronoun condition (when only one of the entities is mentioned in the continuation), equally coherent transitions arise regardless of whether the subject-position pronoun refers to the preceding subject or object. (These predictions are explained in detail in the section entitled "Referential properties of the pronoun-containing clause".) Thus, we predict that *object bias will be relatively stronger in the One-Pronoun condition than the Two-Pronoun condition*. (Or, conversely, that subject bias is stronger in the Two-Pronoun than in the One-Pronoun condition.)

According to the Referential Structure Hypothesis, these effects should hold even when the two clauses are not linked via semantic parallelism. Thus, this diverges from work on parallelism effects on pronoun resolution (e.g. Chambers & Smyth, 1998; Smyth, 1994; Stevenson et al., 1995). To assess this, we chose to use sentences linked by explanation relations indicated by *because*.

In sum, if pronoun interpretation is influenced both by verbs' implicit causality *and* by referential structure, we predict that in addition to subject-position pronouns showing a stronger subject preference with SE (IC1) verbs than with ES (IC2) verbs (main effect of verb type), we also expect a weakening of this subject preference in One-Pronoun compared to the Two-Pronoun conditions (main effect of referential structure). (We do not make any predictions regarding interactions between verb type and referential structure; according to the Referential Structure Hypothesis, we should see effects of referential structure regardless of verb type).

Data analysis

To analyse our data, we used generalised linear mixed logit models, using R (R Core Team, 2019). The dependent variable was the proportion of object interpretations, and models included fixed effects of referential structure (contrast-coded, One-Pronoun = 0.5, Two-Pronoun = −0.5), verb type (contrast-coded, IC1 = −0.5, IC2 = 0.5), and the referential structure x verb type interaction.

In addition, we entered intercepts for subjects and items as random effects into our models, as well as by-subject and by-item random slopes for the effects of IC verb type, referential structure type, and their interaction when justified by model comparison. We started out with fully crossed and fully specified random effects, with by-subject and by-item effects of IC verb type, referential structure type, and their interaction. These random effects were then reduced (starting with by-item effects) by using model comparison, such that only random effects that contributed significantly to the model ($p < 0.05$) were included (Baayen et al., 2008).

Results

Figure 2 shows the percentage of trials on which the subject-position pronoun was interpreted as referring to the preceding *object*, as function of verb type (ES vs. SE) and referential structure (One-Pronoun vs. Two-Pronoun).

As can be seen in Figure 2, referential structure influences pronoun interpretation in the ES (IC2) condition in the direction predicted by the Referential Structure Hypothesis: With ES verbs, the presence of an object-position pronoun in the second clause (Two-Pronoun condition) markedly reduces the percentage of object interpretations (69.6%) relative to configurations without an object-position pronoun (One-Pronoun condition) (92.1%). However, with SE verbs, the Two-Pronoun condition (43.3%) does not differ from the One-Pronoun condition (42.1%) in this regard: Figure 2 suggests that in the SE verb conditions, pronoun interpretation is not sensitive to the One-Pronoun vs. Two-Pronoun manipulation.

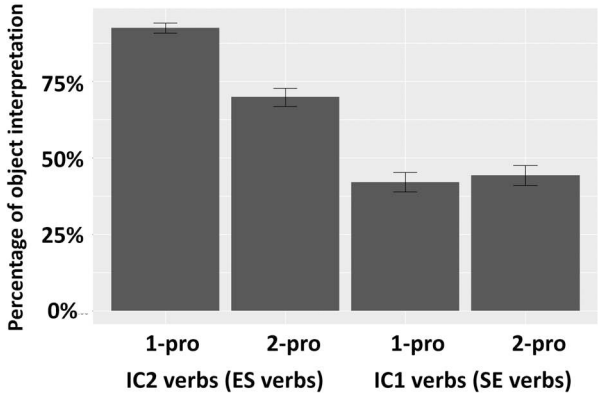

Figure 2. Experiment 1: Percentage of trials where the subject-position pronoun refers to the preceding object. Error bars show ±1 SE.

Visually, Figure 2 also points to implicit causality effects being at play: the proportion of object interpretations is higher with ES (IC2) verbs than SE (IC1) verbs, in both the One-Pronoun and the Two-Pronoun conditions. In the One-Pronoun condition, subject position pronouns are interpreted as referring to object antecedents 92.1% of the time with ES verbs, but only 42.1% of the time with SE verbs. Similarly, in the Two-Pronoun condition, subject position pronouns refer to object antecedents 69.6% of the time with ES verbs, but only 43.3% of the time with SE verbs.

The outcomes of the statistical analysis are shown in Table 2. We find main effects of referential structure, IC verb type and a significant interaction. Further planned comparisons (Table 3) show that the effect of referential structure is significant in the ES conditions, but not in the SE conditions, as expected based on Figure 2. The patterns in the ES condition support the prediction of the Referential Structure Hypothesis that object preference is stronger in the One-Pronoun than in the Two-Pronoun conditions. However, we found no effect of the referential structure in the SE condition.

Discussion

Experiment 1 tested whether pronoun resolution is guided by the referential structure of the pronoun-containing clause, in addition to the implicit causality of the verb in the preceding clause. As regards implicit

Table 2. Experiment 1: results of the glmer model (see Appendix D for the model).

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	0.81	0.22	3.67	<.001	***
Ref.st type	0.85	0.20	4.27	<.0001	***
IC type	2.31	0.19	11.81	<.0001	***
Ref.st type:IC type	1.96	0.38	5.16	<.0001	***

Table 3. Experiment 1: planned comparisons (see Appendix D for the models).

		Estimate	Std. Error	z value	Pr(>|z|)	
SE-verbs	(Intercept)	−0.37	0.30	−1.23	.21	
	Ref.st type	−0.17	0.22	−0.77	.43	
ES-verbs	(Intercept)	2.04	0.28	7.04	<.0001	***
	Ref.st type	1.99	0.32	6.07	<.0001	***

causality, our results replicate earlier results: subject-position pronouns were more likely to be interpreted as referring to objects after ES (IC2) verbs than after SE (IC1) verbs. This is in line with prior work (e.g. Caramazza et al., 1977; Ferstl et al., 2011; Hartshorne & Snedeker, 2013; McKoon et al., 1993) and also serves as a "sanity check" to show that the expected patterns replicate with our novel picture-writing task and with nonce verbs in the second clause.[2]

As regards referential structure, we find significant effects with Experiencer-Stimulus (ES) verbs: the One-Pronoun configuration shows a stronger object bias than the Two-Pronoun configuration. This finding supports the Referential Structure Hypothesis, which predicts that the interpretation of subject-position pronouns in One-Pronoun conditions exhibits a stronger object preference (and a weaker subject preference) than in Two-Pronoun conditions.

However, with Stimulus-Experiencer (SE) verbs, contrary to our predictions, we fail to see a stronger object preference in the One-Pronoun condition relative to the Two-Pronoun condition. There are two points we would like to highlight regarding this result.

First, the null result with SE verbs coupled with a significant Referential Structure effect with ES verbs indicates that referential structure effects cannot be reduced simply to effects of surface-level syntactic parallelism, because the One- and Two-Pronoun conditions have the same surface syntactic structure regardless of verb type.

Second, the null result with SE verbs may stem from a particular semantic property of SE verbs coinciding with a potential confound between the One-Pronoun and Two-Pronoun conditions. To see why, let us first note that in the Two-Pronoun condition, the pronoun-containing clauses contained transitive (nonce) verbs (e.g. *she frobbed her*), whereas in the One-Pronoun condition, these clauses contain intransitive (nonce) verbs (e.g. *she frobbed*). This transitivity manipulation allowed us to construct One-Pronoun vs. Two-Pronoun configurations, but it is also associated with semantic differences: prior work shows that subjects of intransitive verbs are less semantically agentive than subjects of transitive verbs (e.g. Dixon, 1979; Hopper & Thompson, 1980). This means that the referents of the subject pronouns in the (intransitive) One-Pronoun conditions are likely to be regarded as less agentive and less volitional – and thus also lower-ranked on the thematic hierarchy – than the referents of the subject pronouns in the Two-pronoun conditions. Second, as mentioned above, subjects of ES and SE verbs also differ in prominence: the Stimulus subjects of SE verbs are widely viewed as less thematically prominent than the Experiencer subjects of ES verbs (see e.g. Belletti & Rizzi, 1988; Bresnan & Kanerva, 1989; Foley, 2005; Foley & Van Valin, 1984; Giorgi, 1984; Grimshaw, 1990; Speas, 1990; Van Valin, 2001; Van Valin & LaPolla, 1997).

Put together, these two considerations (intransitive subjects being less prominent and SE subjects being less prominent) mean that sentences in the *SE (IC1) + One-Pronoun* condition are likely to be interpreted as having relatively non-prominent subjects in both clauses. This brings up a potential alternative explanation for the results in the SE condition of Experiment 1: If pronoun interpretation is guided by thematic role prominence such that pronouns with a relatively lower prominence thematic role prefer antecedents with a correspondingly lower prominence thematic role, this could result in the subject preference in the *SE (IC1) + One-Pronoun* condition being boosted (and the object preference being lowered) – which would explain why there is no referential structure effect with SE verbs (IC1). To address this concern and to test whether we find support for the Referential Structure Hypothesis when the thematic roles in the One-Pronoun and Two-Pronoun configurations are held constant, Experiment 2 used transitive verbs in both clauses.

Experiment 2

The results of Experiment 1 provide evidence that the referential structure of the pronoun-containing clause guides pronoun interpretation in contexts with Experiencer-Stimulus (ES) verbs, but show no effects of referential structure with Stimulus-Experiencer (SE) verbs. As discussed above, this unexpected lack of across-the-board referential structure effects may be due to concerns related to verb transitivity and associated semantic properties. To test the Referential Structure Hypothesis more directly, in Experiment 2 the nonce verbs in the pronoun-containing clause were consistently transitive. The Two-Pronoun conditions were as in Experiment 1 (e.g. *she frobbed her*) and the One-Pronoun conditions now consist of transitive sentences with a new name in object position (e.g. *she frobbed Kate*). We will refer to these as Pronoun + Name conditions.

Method

Participants

Forty-eight native English-speaking adults (age 18 or over), none of whom had participated in Experiment 1, participated in this study. Participant recruitment was the same as in Experiment 1. Four participants were excluded because they were not self-identified U.S.-born native speakers of English, gave random answers

throughout the task, or did not pass the minimum catch trial threshold. (This experiment had 11 catch trials; participants had to get at least 8 correct).[3] We excluded an additional four participants to balance the number of participants on each list. A total of forty participants (20 self-identified as female, 19 as male, 1 other) were included in the final analyses.

Materials and design

The design and materials of Experiment 2 were the same as Experiment 1, except that we now manipulated referential structure without introducing an (in)transitivity confound, as shown in Table 4. The critical pronoun-containing clause is now transitive in both the Two-Pronoun condition (e.g. *she kreeged her*) and the Pronoun + Name condition (e.g. *She kreeged Jocelyn*). In addition to the 24 targets, Experiment 2 included 36 fillers (including 11 that functioned as catch trials).

Procedure

The task and the procedure were the same as in Experiment 1. All target items now involve two characters (Figure 3), because all nonce verbs in the pronoun-containing clause denote transitive events. The displays that participants saw in the Two-Pronoun conditions were the same as in Experiment 1 (Figure 1a). An example display for the Pronoun + Name condition is provided in Figure 3.

Table 4. An example item from Experiment 2.

Verb Type	Referential structure	Examples
IC1 (SE)	Two-Pronoun	Stacy discouraged Hanna because she kreeged her.
	Pronoun + Name	Stacy discouraged Hanna because she kreeged Jocelyn.
IC2 (ES)	Two-Pronoun	Stacy believed Hanna because she kreeged her.
	Pronoun + Name	Stacy believed Hanna because she kreeged Jocelyn.

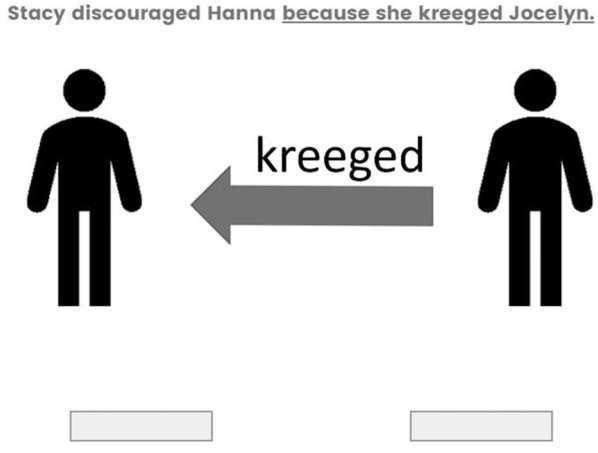

Figure 3. Experiment 2: Sample Pronoun + Name item.

Data processing

We coded the data in the same way as in Experiment 1. Of all the data, 1.1% were excluded (marked as N/A) because the subject-position pronoun or object-position pronoun was not interpreted as mentioning the subject or object of the preceding clause. (As in Experiment 1, these are presumably errors.) In total, 98.9% of the data were submitted for statistical analysis.

Predictions

The predictions are the same as for Experiment 1.

Data analysis

Statistical analyses were conducted the same way as Experiment 1.

Results

Figure 4 presents the percentage of trials on which the subject-position pronoun is interpreted as referring to the preceding *object*. The dependent variable was again the proportion of object interpretations, and models included fixed effects of referential structure (contrast-coded, Pronoun + Name = 0.5, Two-Pronoun = −0.5), verb type (contrast-coded, IC1 = −0.5, IC2 = 0.5), and the referential structure x verb type interaction. Random effects were determined as in Experiment 1.

In Experiment 2, we find a main effect of IC verb type in the expected direction (more object interpretations with ES/IC2 verbs), a main effect of referential structure

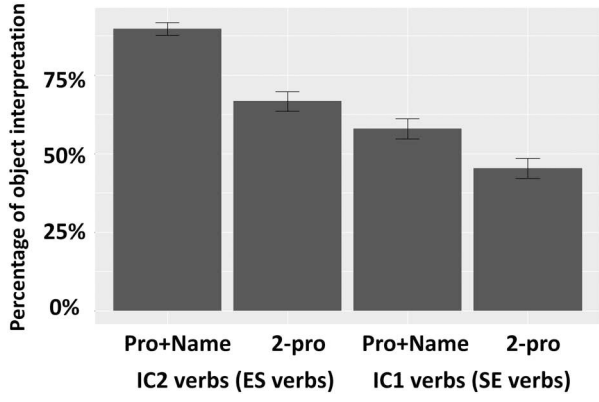

Figure 4. Experiment 2: Percentage of trials where the subject-position pronoun refers to the preceding object. Error bars show ±1 SE.

Table 5. Experiment 2: results of the glmer model (see Appendix D for the model).

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	0.89	0.21	4.17	<.0001	***
Ref.st type	1.16	0.19	6.02	<.0001	***
IC type	1.66	0.18	9.16	<.0001	***
Ref.st type:IC type	1.07	0.35	3.02	.002	**

Table 6. Experiment 2: planned comparisons (see Appendix D for the models).

		Estimate	Std. Error	z value	Pr(>|z|)	
SE-verbs	(Intercept)	0.09	0.26	0.36	.71	
	Ref.st type	0.66	0.24	2.71	.006	**
ES-verbs	(Intercept)	3.59	1.19	3.00	.002	**
	Ref.st type	5.37	2.29	2.34	.01	*

(more object interpretations with Pronoun + Name conditions than Two-Pronoun conditions), and an interaction between verb type and referential structure (stronger effects of referential structure with ES than SE verbs). See Table 5 for statistical details. Importantly, planned comparisons confirm that there are more object interpretations in the Pronoun + Name configurations than in the Two-Pronoun configurations with *both* SE verbs and ES verbs (Table 6). Thus, although we obtain an interaction, Experiment 2 – in contrast to Experiment 1 – reveals significant effects of referential structure with both verb types.

Discussion

The results of Experiment 2 echo Experiment 1 in showing that both referential structure and IC verb bias guide pronoun interpretation. In a set-up where we eliminate potential confounds from transitivity, we now find significant referential structure effects with both SE and ES verbs: the preference for object antecedents is stronger in the Pronoun + Name configuration (where only one of the referents from the preceding clause is mentioned) than in the Two-Pronoun configuration (where both of the referents from the preceding clause are mentioned), with both verb types. This finding supports the Referential Structure Hypothesis, which posits that the ultimate interpretation of pronouns is influenced by the referential structure of the entire pronoun-containing clause, i.e. whether the other referent from the preceding clause is also mentioned in that clause. (As mentioned in the introduction, this finding is not straightforwardly predicted by an account relying purely on syntactic parallelism, because the syntactic structures and grammatical roles in the Pronoun + Name and Two-Pronoun conditions do not differ.)

We also find significant effects of IC verb type in the expected direction (more object interpretations with ES/IC2 verbs than SE/IC1 verbs). This confirms that the expected IC effects replicate with our task. Furthermore, the IC effects reveal that participants are attending to fine-grained verb semantics (the differences between IC1 and IC2 verbs) and are not simply engaging in shallow, one-to-one mapping of stick figures to pronouns without attending to the linguistic properties of the sentences.

We also find an interaction between IC verb type and referential structure: the referential structure effect is present with both ES and SE verbs, but stronger with ES verbs than with SE verbs. This verb-type difference, while entirely compatible with the Referential Structure Hypothesis, is not directly predicted by it. We speculate that it may stem from differences in the syntactic-thematic alignment patterns that we mentioned in the introduction. Recall that with ES verbs, thematic and syntactic prominence are aligned (the high-ranked Experiencer role is also syntactically high-ranked, as the subject), while with SE verbs, they are misaligned (the highest-ranked thematic role is not in the highest-ranked syntactic position). We return to this in the General Discussion section.

As a whole, Experiment 2 provides evidence in favour of the Referential Structure Hypothesis. However, both Experiments 1 and 2 used psych verbs (ES/SE verbs), with Experiencer and Stimulus thematic roles. This raises the question of whether referential structure effects also generalise to other verb classes. To assess the robustness of referential structure effects, Experiment 3 tests Agent-Patient and Agent-Evocator verbs. These verb classes have clearly agentive subjects and describe the actions of the event participants, in contrast to ES/SE psych verbs that describe experiencers' mental states.

Experiment 3

Experiment 3 tests whether referential structure effects extend to action verbs with different thematic roles than the psych verbs used in Experiments 1 and 2. Given that the Referential Structure Hypothesis makes no reference to verb semantics, it predicts that referential structure effects should occur with all verb types (other things being equal). Furthermore, given that the verb types tested in Experiment 3 – Agent-Patient (AP, IC1) and Agent-Evocator (AE, IC2) – do not exhibit misalignment between the mapping of thematic roles to syntactic positions (the highest-ranked thematic

Table 7. An example item from Experiment 3.

Verb Type	Referential structure	Examples
IC1 (AP)	Two-Pronoun	Stacy followed Hanna because she kreeged her.
	Pronoun + Name	Stacy followed Hanna because she kreeged Jocelyn.
IC2 (AE)	Two-Pronoun	Stacy welcomed Hanna because she kreeged her.
	Pronoun + Name	Stacy welcomed Hanna because she kreeged Jocelyn.

role is also in the highest-ranked syntactic position), the results can also help shed light on the ES/SE differences we found in Experiment 2.

Method

Participants

Sixty participants were included in the final analyses (32 self-identified as female, 28 as male), out of 71 native English-speaking adults (18 or older) who took part. The recruitment and exclusion criteria were the same as Experiments 2. We excluded three participants because they were not born in the US and/or were not native speakers of English. Five participants were excluded because they consistently gave random answers or did not pass the threshold for catch trials (at least 8/11 correct). Three more participants were excluded to balance the number of participants on each list.

Materials and design

The materials and design are the same as Experiment 2, except that the first clause now used Agent-Patient (AP) verbs (which have an IC1 bias) and Agent-Evocator (AE) verbs (which have an IC2 bias), instead of ES/SE verbs.[4] Examples are in Table 7. Twenty-four AP and AE verbs were selected based on the norms collected by Hartshorne and Snedeker (2013) and Ferstl et al. (2011) to ensure that the AP verbs are indeed subject-biased (mean subject bias = 67.7%, SD = 9.16) and that the AE verbs are object-biased (mean object bias = 72.1%, SD = 5.53). Other than this verb change, Experiment 3 used the same 24 targets and the same 36 fillers (including 11 that function as catch trials) as Experiment 2.

Procedure

The procedure was the same as Experiments 1 and 2, using Qualtrics (Provo, UT, 2020). Examples of the displays are in Figure 5.

Data processing

Experiment 3 was analysed like Experiments 1 and 2. Of all the data, 0.3% were removed (marked as N/A) because the subject-position pronoun or object-position pronoun was not interpreted as referring to either the preceding subject or object. In total, 99.7% of the data were submitted for analysis.

Predictions

If the Referential Structure Hypothesis generalises across verb classes, the predictions are the same as for Experiments 1 and 2.

Data analysis

Statistical analyses were conducted the same way as in Experiments 1 and 2.

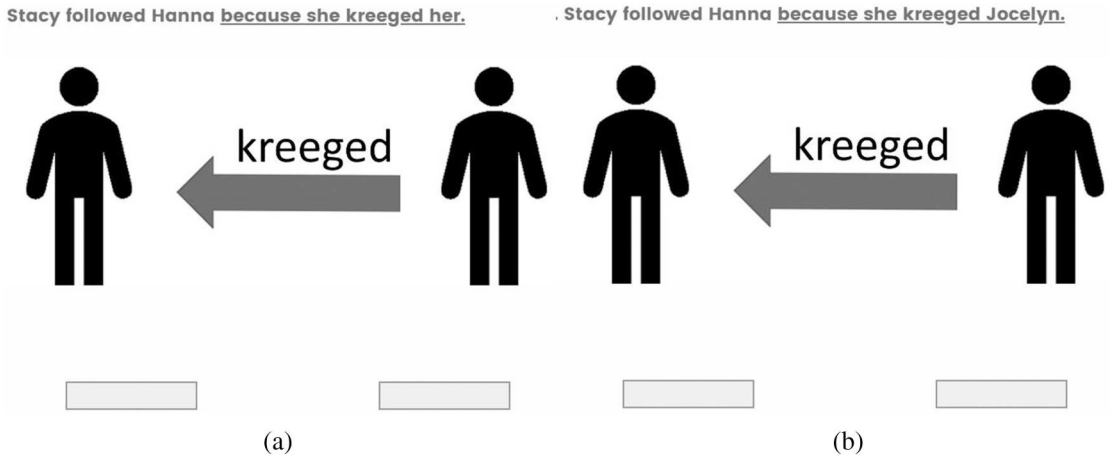

Figure 5. (a) Experiment 3: Sample Two-Pronoun item. (b) Experiment 3: Sample One-Pronoun item.

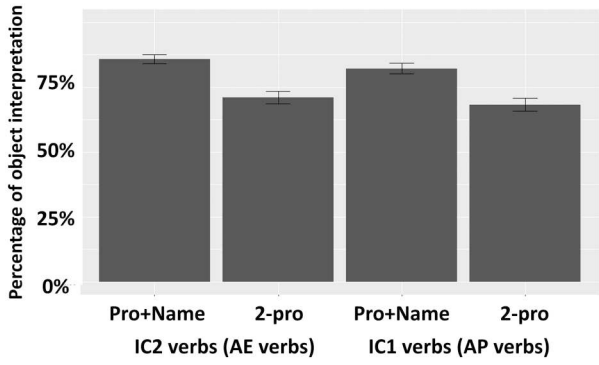

Figure 6. Experiment 3: Percentage of trials where the subject-position pronoun refers to the preceding object. Error bars show ±1 SE.

Results

Figure 6 shows the percentage of trials where the subject-position pronoun is interpreted as referring to the *object* of the preceding clause. Just as in Experiment 1 and Experiment 2, the dependent variable was the proportion of object interpretations, and models included fixed effects of referential structure (contrast-coded, Pronoun + Name = 0.5, Two-Pronoun = −0.5), verb type (contrast-coded, IC1 = −0.5, IC2 = 0.5), and the referential structure x verb type interaction.

Experiment 3 reveals main effects of both referential structure and IC verb type, in the expected directions, and no interaction (Table 8). Thus, Experiment 3 replicates the expected IC effects (more object interpretations with object-biased AE verbs than with subject-biased AP verbs). As regards referential structure, Figure 6 shows that subject-position pronouns exhibit a stronger object preference in Pronoun + Name conditions than in Two-Pronoun conditions. The absence of an interaction shows that this preference is equally strong with AE verbs and AP verbs. This "across-the-board" increase in object choices in the Pronoun + Name condition relative to the Two-Pronoun condition is exactly in line with the predictions of the Referential Structure Hypothesis.

Discussion

Experiment 3 replicated and extended the findings of Experiments 1 and 2, which tested Experiencer-Stimulus (ES) and Stimulus-Experiencer (SE) verbs, to agentive verb classes, namely Agent-Patient (AP) and Agent-Evocator (AE) verbs. Crucially, we again find effects of referential structure: The preceding object was more likely to be interpreted as the antecedent of a subsequent subject-position pronoun in the Pronoun + Name than in the Two-Pronoun configuration. These results support the Referential Structure Hypothesis which states that a subject-position pronoun is more likely to be interpreted as referring to the preceding object when the pronoun-containing clause only mentions one referent from the prior sentence. Combined with Experiments 1 and 2, Experiment 3 demonstrates the robustness of the referential structure effect: It occurs with verbs that denote agentive, voluntary actions (AP/AE verbs) in addition to verbs that denote mental states (SE/ES verbs).

Furthermore, in line with prior work on implicit causality, Experiment 3 shows that subject-position pronouns are more likely to be interpreted as referring to the preceding subject with AP (IC1) verbs than AE (IC2) verbs. Thus, Experiment 3 extends the results of Experiments 1 and 2 by further corroborating effects of implicit causality and shows that these effects can be reliably detected with our experimental paradigm and in the presence of nonce verbs.

Overall, the results of Experiment 3 confirm that, ultimately, pronoun resolution is guided not only by the IC information in the preceding clause but also by the referential structure of the entire pronoun-containing clause, indicating that the referential structure effects are strong enough to be observed with verbs of different semantic classes.

Table 8. Experiment 3: results of the glmer model (see Appendix D for the model).

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	2.31	0.38	6.08	<.0001	***
Ref.st type	1.69	0.37	4.58	<.0001	***
IC type	0.39	0.18	2.16	.03	*
Ref.st type:IC type	0.32	0.36	0.89	.37	

Comparing the three experiments

So far, we have seen evidence from three experiments supporting the Referential Structure Hypothesis – i.e. the idea that pronoun interpretation is ultimately sensitive to whether a clause mentions only one or both of the referents from the preceding clause. To the best of our knowledge, our work is the first experimental demonstration of this effect.

Orthogonal to this key point, it is interesting to note that, at least visually, a comparison of Figures 2, 4, and 6 suggests that the rate of object interpretations is higher in Experiment 3 (with agentive action verbs) than in Experiments 1 and 2 (with ES/SE psych verbs). Indeed, statistical analyses (glmer) comparing the proportion of object interpretations with AP (IC1) verbs (Experiment 3) relative to SE (IC1) verbs (Experiments 1 and 2) confirm that AP verbs elicit higher rates of object interpretations than SE verbs (Exp3 vs. Exp1: beta = −2.22, SE = 0.46, z = −4.74, $p < .001$; Exp3 vs.

Exp2: beta = −1.74, SE = 0.44, z = −3.85, $p < .001$). Thus, in our experiments, AP (IC1) verbs elicit higher rates of object interpretations than SE (IC1) verbs. This is the case even though based on corpus norms, the subject bias strengths of the AP verbs (mean subject bias = 67.7%, SD = 9.16) and SE verbs (mean subject bias = 67.4%, SD = 13.6) do not differ significantly (t(46) = 0.35, $p > .7$). In fact, a similar asymmetry was reported by Stevenson et al. (1994), who used a variety of verb classes and connectives in fragment-completion studies (e.g. *Ken admired Geoff ./and/because/so he* …).[5]

Crucially, this difference in object bias strength between AP and SE verbs is orthogonal to the main claims we are making, especially since we make no predictions about (and nothing hinges on) absolute differences between verb types' object bias strength. For us, what matters is that both Experiments 2 and 3 show effects of verb bias as well as effects of referential structure.

General discussion

Despite the large body of psycholinguistic work on reference resolution, most prior investigations have focused on the production and comprehension of pronouns in subject position, with little attention given to the potential presence or consequences of pronouns occurring later in the same sentence (with the notable exception of work on parallelism). In the present paper, we highlight the importance of looking beyond pronouns in subject position: our findings show that comprehenders' ultimate interpretation of subject-position pronouns is modulated by the referential structure of the entire pronoun-containing clause. We show that interpretation of subject-position pronouns is also guided by factors that occur *after* these pronouns, even in contexts that do not involve semantic parallelism.

Building on ideas adapted from Centering Theory (e.g. Grosz et al., 1995; Walker et al., 1998), we propose and test the **Referential Structure Hypothesis**, according to which subject pronouns in clauses that only mention one of the referents from the preceding clause are more likely to be interpreted as referring to a preceding object than subject pronouns in clauses that mention both referents from the preceding clause (e.g. *she verbed Kate* vs. *she verbed her*).

To look for effects of referential structure, we compared (i) configurations where the second clause only mentions one of the two entities from the preceding clause with a subject-position pronoun (e.g. *because she kreeged* vs. *because she kreeged Jocelyn*) to (ii) configurations where the pronoun-containing clause mentions both entities from the preceding clause with pronouns (e.g. *she kreeged her*). We conducted three experiments using both non-agentive and agentive verbs in the preceding clause (non-agentive IC1 and IC2 verbs [SE/ES verbs] in Experiments 1 and 2, agentive IC1 and IC2 verbs [AP/AE verbs] in Experiment 3), to assess the robustness of referential structure effects in different contexts.

Referential structure effects

In all three experiments, we find overall effects of referential structure on pronoun interpretation: a pronoun in the subject position is more likely to be interpreted as referring back to the preceding subject in configurations with two pronouns (e.g. *she verbed her*) than in configurations with only one pronoun (e.g. *she verbed, she verbed Kate*).

More specifically, in Experiment 2 and Experiment 3, the referential structure effects are observed with both IC1 and IC2 verbs, while in Experiment 1, these effects are only observed with IC2 verbs (not IC1 verbs). We suggest that the lack of referential structure effects with IC1 verbs in Experiment 1 likely stems from the fact that Experiment 1 compared transitive and intransitive clauses, whose subjects differ in thematic prominence. Once this potential confound was addressed in Experiments 2 and 3, we found referential structure effects with both IC1 and IC2 verbs. These results were obtained with clauses that have the same surface syntactic structure and differ only in their referential properties, suggesting that a purely syntax-based account cannot straightforwardly capture our findings. In other words, regardless of whether the verb has a baseline bias towards the preceding subject (IC1) or the preceding object (IC2) in explanation contexts like the ones we tested, the likelihood of a subject-position pronoun being interpreted as referring to the preceding subject vs. object is influenced by the presence/absence of another pronoun in the same clause.

These results are compatible with a fundamental insight based on Centering Theory – namely that the referential properties of clauses contribute to discourse coherence and that the interpretation of pronouns is driven by a bias to maximise discourse coherence. (However, as noted in the introduction, our Referential Structure Theory is not intended to be viewed as an implementation or test of Centering Theory, and differs in several of its assumptions.)

Let us first consider two-pronoun sentences. The basic idea is that, in a context where one clause (e.g. *Lisa verbed Mary because* …) is followed by a second clause with both subject- and object-position pronouns (e.g. *she verbed her*), interpreting the subject-position pronoun as referring to the previously lower-salience

referent (preceding object) and the object-position pronoun as referring to the previously higher-salience referent (preceding subject) yields a non-ideal transition. This is because a previously high-salience referent (the subject *Lisa*) loses out to a previously low-salience referent (the object *Mary*) in terms of their subject and object roles in the second clause. In contrast, interpreting the subject-position pronoun as referring to the preceding subject and the object-position pronoun as referring to the preceding object yields a highly coherent transition. In other words, an interpretation that demotes the preceding subject to a less prominent (object) position and promotes the preceding object to a more prominent (subject) position is dispreferred.

Now, let us consider One-Pronoun cases. If the second clause only contains a subject-position pronoun (e.g. *she verbed, she verbed Kate*), then the transition is equally coherent regardless of whether the pronoun is interpreted as referring to the preceding subject or object. A coherent transition obtains regardless of whether we (a) *promote* a previously lower-salience object to the more prominent subject position or (b) *maintain* the already higher-salience preceding subject in the prominent subject position. This is because in neither configuration does a previously high-salience referent get outranked by a previously lower-salience referent.

Moreover, these findings regarding pronoun interpretation are compatible with Kaiser's (2009) production results – where referential structure and preceding verb type were not explicitly manipulated – which suggest that if a preceding object is promoted to the subject position of a following clause (by means of a subject-position pronoun), the preceding subject is less likely to be mentioned later in the same clause.

A broader perspective on referential structure

Effects of referential structure may be related to broader, domain-general perceptual notions of similarity, parallelism and symmetry. Although the studies we report in this paper were not designed to directly tap into this possibility, it may well be that an abstract cognitive notion of parallelism of some kind is at play. It is well-known that human perception in various non-linguistic domains (e.g. vision, audition, tactile/haptic processing) is guided in deep ways by abstract notions related to similarity, parallelism and symmetry (e.g. gestalt notions and related work, see e.g. Palmer (1999) Chapter 6 on vision). Research on these kinds of effects has led to fundamental insights about the domain-general mechanisms that guide human perception. Given that notions of similarity and parallelism have been argued to guide human perception and cognition in an abstract, domain-general way, then – depending on how we assume language processing to interface with other aspects of human cognitive processing – they may play a role in language processing as well.

Thus, it could be that the Reference Structure Hypothesis is a specific version of (or related to) a more general human perceptual preference for symmetry or parallelism. In our case, the idea would be that matching numbers of referents and pronouns across two clauses in the Two-Pronoun condition guide resolution of referential dependencies, even when the semantic relation between the two clauses does *not* involve semantic parallelism. If this idea is on the right track, it would provide new evidence that aspects of discourse-level processing (here, pronoun interpretation) are influenced by domain-general perceptual mechanisms. This possibility opens the door for new investigations into broader questions about domain generality, mental representations and the relation between linguistic and non-linguistic representations. These questions go beyond the scope of the current paper, which instead sought to take the necessary first step of testing *whether* there are indeed referential structure effects in pronoun resolution – a question that had been left largely unanswered by prior work.

Implications for past and future experimental investigations

As a whole, the results of all three comprehension experiments indicate that referential properties of the entire pronoun-containing clause are utilised during pronoun resolution. These findings have empirical and methodological implications for existing work using the sentence-completion/sentence-continuation paradigm. In this widely-used paradigm, participants are typically given a clause (e.g. *Lisa impressed Tom*), with or without a connective and/or a pronoun prompt at the end (e.g. *Lisa impressed Tom because he/she*), and asked to write a completion. It is currently standard practice to focus only on the subject of the second clause (e.g. Arnold, 1998; Colonna et al., 2012, 2014; Cowles, 2007; Hartshorne & Snedeker, 2013; Johnson & Arnold, 2021; Kaiser, 2011a, 2019; Koornneef & Sanders, 2013; Rohde et al., 2006). The question of whether the other referent from the initial clause is re-mentioned in a later non-subject position in the continuation is, to the best of our knowledge, typically not systematically investigated. However, our results suggest that the rest of the continuation should also be analysed, and that focusing only on the subject position may lead researchers to overlook meaningful differences between conditions (in terms of whether the other referent is or is not mentioned).

On a very broad level, abstracting away from the specific contexts and coherence relations that we tested, we think it is worthwhile to explore the possibility that looking only at the referential properties of the subject brings with it potential risks of incorrectly concluding that (i) two discourse configurations pattern alike, when their subsequent referential structure could potentially reveal significant differences, or conversely, incorrectly concluding that (ii) a referent is more likely to be mentioned or referred to again in one configuration than in another, when actually it is mentioned in equally both but just in a different position in the sentence. We suggest that considering a sentence's referential structure beyond the subject can be important for both empirical and theoretical reasons.

It is worth emphasising that the present work focused on participants' end-of-sentence interpretations, because our key aim in this initial set of studies was to establish whether participants' final pronoun interpretation patterns are influenced by referential structure. We view this work as an initial step that lays the groundwork necessary for future work on real-time, incremental processing. Referential structure effects pose an interesting challenge for incremental processing: If information after the subject pronoun (e.g. the presence of a second pronoun) guides the interpretation that participants ultimately opt for, then this means that, during real-time processing, participants may interpret a pronoun in a way that will subsequently need to be re-analysed, in light of subsequent information. Although the present studies did not aim to address these online processing questions, the results of Experiments 1, 2 and 3 provide (to the best of our knowledge) the first systematic experimental evidence of referential structure effects on end-of-sentence pronoun interpretation – thus providing a crucial foundation for future processing studies using time-sensitive methods such as reaction times, eye-tracking and ERP (see also Brilmayer & Schumacher, 2021 for recent ERP work on pronoun processing).

Implicit causality effects

Another aim of our study was to test whether well-known implicit causality (IC) effects are replicable with our novel method (the picture-writing task). We manipulated verb semantics in the preceding clause by using different classes of IC verbs, and compared IC verbs known to have a subject bias (IC1 verbs) to IC verbs known to have an object bias (IC2 verbs). Experiment 1 and Experiment 2 used non-agentive IC1 and IC2 verbs with Stimulus and Experiencer arguments. Experiment 3 used IC1 and IC2 verbs with agentive subjects. In all three experiments, a significant effect of verb bias was observed: Subject-position pronouns consistently showed a stronger object preference with IC2 verbs than IC1 verbs. These results replicate prior work (e.g. Bott & Solstad, 2014; Caramazza et al., 1977; Ferstl et al., 2011; Hartshorne et al., 2015; Hartshorne & Snedeker, 2013; McKoon et al., 1993) and thus confirm that our picture-writing task, even with nonce verbs in the second clause, yields meaningful data regarding pronoun interpretation. Thus, our study corroborates, using a novel task, previous findings that verb cues encountered before the pronoun are utilised for pronoun resolution.

Thematic role effects

In our study, we found an unexpected interaction between verb type and referential structure with SE (IC1) / ES (IC2) verbs (Experiments 1,2), but not with AP (IC1) / AE (IC2) verbs (Experiment 3). While sentences with SE (IC1) verbs were less susceptible to referential structure effects than those with ES (IC2) verbs, sentences with AP (IC1) and AE (IC2) verbs were equally sensitive to referential structure effects. Although this asymmetry is not central to or problematic for our claims (given that both SE and ES verbs still showed referential structure effects), let us nevertheless briefly consider what could be causing it.

As noted in the introduction, the IC differences exhibited SE/ES verbs are essentially "confounded" with syntactic-thematic (mis)alignment, while this is not the case with AP/AE verbs. SE verbs involve a misalignment (the thematically highest-ranked argument, Experiencer, is not in the syntactically highest-ranked position), whereas ES verbs are aligned (the high-ranked Experiencer is in subject position). Thus, thematic prominence and syntactic prominence are misaligned. In contrast, both AP and AE verbs are aligned, as the high-ranked Agent is in subject position. Thus, thematic prominence and syntactic prominence are aligned.

First, let us consider how this might yield relatively weaker referential structure effects with SE verbs than ES verbs. Experiencers are known to be inherently more salient than Stimuli (as reflected by Experiencers being higher ranked on the thematic hierarchy). This is often attributed to Experiencers being animate and sentient, while Stimuli are not necessarily animate or sentient (e.g. Fedriani, 2014; Verhoeven, 2009, 2014). Recall that we suggest referential structure effects are related to a dispreference for transitions where a previously higher-salience referent (subject of the preceding clause) is demoted to a lower-salience position (object of the subsequent clause) while a lower-salience referent (preceding object) is simultaneously promoted.

Considering that an Experiencer is more salient than a Stimulus, there could be differences in the consequences of demoting the subject in clauses with SE verbs (subject = Stimulus) vs. clauses with ES verbs (subject = Experiencer). Specifically, demoting a salient Experiencer referent could yield an even less coherent transition than demoting a less salient Stimulus referent. This would result in ES verbs exhibiting greater sensitivity to referential structure effects that SE verbs – which is indeed what we found.

In contrast, the thematic role of the subject does not differ between AP (IC1) verbs and AE (IC2) verbs – the subject is an Agent in both cases. In other words, there is no syntactic-thematic role misalignment. Accordingly, we do not expect any asymmetry in the strength of the referential structure effects. This matches what we found in Experiment 3.

Thus, we suggest that the verb type x referential structure interaction that we observe with SE/ES verbs in particular may stem from differences in the alignment/mapping between syntactic prominence and thematic prominence (on (mis)alignment more generally, see also Fuchs & Schumacher, 2020; Schumacher et al., 2016, 2017, as well as Do & Kaiser, 2022 and Ferreira, 1994 for production-based work). This fits with prior claims that the salience of discourse entities is influenced by thematic roles as well as by grammatical roles.

Conclusions

The three experiments reported here provide novel evidence for the Referential Structure Hypothesis, which posits that whether subject-position third-person pronouns are interpreted as referring to the preceding subject or object depends on whether or not the clause contains another third-person pronoun referring to the other referent in the preceding clause. Our results show that in addition to being influenced by information that comes before the pronoun (verb semantics), pronoun interpretation is also sensitive to the referential properties of the entire pronoun-containing clause. We show that subject-position pronouns have a stronger object preference when only one of the antecedents from the prior clause is mentioned, compared to when both the antecedents are mentioned. We explain this effect in terms of a bias to maximise discourse coherence and show that it is separate from effects of verb semantics, and cannot be reduced to a semantic parallelism effect or a pure syntactic parallelism effect. As a whole, the studies reported here provide novel evidence that models of pronoun resolution should be extended to include referential structure.

Notes

1. Previous papers on IC verb effects that used nonce words (e.g. Hartshorne & Snedeker, 2013; Hartshorne et al., 2015) show that meaningful results about pronoun resolution patterns can be successfully obtained with designs using nonce words (see also Burnsky et al., 2022). In addition, earlier results from Kaiser (2009) using a paradigm without nonce words are compatible with our findings. These prior studies suggest that the results reported in the present paper are not artifact stemming from our use of nonce verbs.
2. In this experiment (and the two following ones as well), the absolute proportion of object interpretations is somewhat higher than one might expect given the existing norms (Ferstl et al., 2011; Hartshorne & Snedeker, 2013). A similar increase in object interpretations was also observed by Hartshorne et al. (2015). Crucially, this is not problematic for the main claims we are making, since we are interested in relative differences between IC1 and IC2 verbs (which clearly obtain in our results), not in absolute numbers.
3. The number of catch trials in Experiment 2 is higher than in Experiment 1, because Experiment 2 included transitive as well as intransitive sentences as unambiguous catch trials. This was done to ensure participants saw a mix of transitive and intransitive sentences over the course of Experiment 2 (where all targets were transitive), to keep it similar to Experiment 1 (which had both transitive and intransitive targets).
4. As mentioned earlier, the Evocator role is equivalent to Patient, semantically. The term Evocator is used simply to distinguish agent-patient verbs that elicit IC1 bias (Agent-Patient) from those that elicit IC2 bias (Agent-Evocator, (e.g. Au, 1986; Ferstl et al., 2011; Rudolph & Förster-ling, 1997). Thus, for our purposes, all verbs in Experiment 3 have an agent argument and a patient argument.
5. In Stevenson et al. (1994), a subject-position pronoun preferred a Patient antecedent over an Agent antecedent with AP verbs, but no overall preference for a certain semantic role was found with SE/ES verbs (it varied depending on connective type). Stevenson et al. argued that the patient preference "reflects the salience of the consequences of the described events. Since experiencer-stimulus sentences describe states rather than events, they are not subject to comparable focusing in the constructed representation" (Stevenson et al., 1994, p. 540).

Acknowledgements

Thanks to Ian Rigby and Jesse Storbeck for help with creating the experimental stimuli. We would like to thank audiences at the 3rd California Meeting on Psycholinguistics (CAMP) 2019 (University of California, Santa Cruz, USA), the CUNY Conference on Human Sentence Processing 2020 (University of Massachusetts-Amherst, USA), the 42nd Annual Meeting of the Cognitive Science Society (COGSCI) 2020, the CUNY Conference on Human Sentence Processing 2021 (University of Pennsylvania, USA), the 2021 Winter LSK Young Scholar Symposium (Seoul, Korea), the 4th California Meeting on Psycholinguistics (CAMP) 2021 (University of California, Irvine, USA), and the Human Sentence Processing (HSP) conference 2022 (University

of California, Santa Cruz, USA), in which earlier versions of some of this research were presented. An earlier version of some of the results from Experiment 1 and Experiment 2 was included in the *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society* (Song & Kaiser, 2020b).

Disclosure statement

No potential conflict of interest was reported by the author(s).

Ethics statement

The studies involving human participants were reviewed and approved by the University of Southern California Institutional Review Board. Written informed consent was waived for this study.

ORCID

Jina Song http://orcid.org/0000-0002-5560-1271
Elsi Kaiser http://orcid.org/0000-0003-3594-4127

References

Ariel, M. (1990). *Assessing noun-phrase antecedents*. Routledge.
Ariel, M. (1994). Interpreting anaphoric expressions: A cognitive versus a pragmatic approach. *Journal of Linguistics*, *30*(1), 3–42. https://doi.org/10.1017/S0022226700016170
Arnold, J. (1998). *Reference form and discourse patterns* [Unpublished doctoral dissertation]. Stanford University.
Arnold, J. E., Kaiser, E., Kahn, J. M., & Kim, L. K. (2013). Information structure: Linguistic, cognitive, and processing approaches. *Wiley Interdisciplinary Reviews: Cognitive Science*, *4*(4), 403–413. https://doi.org/10.1002/wcs.1234
Au, T. K. (1986). A verb is worth a thousand words: The causes and consequences of interpersonal events implicit in language. *Journal of Memory and Language*, *25*(1), 104–122. https://doi.org/10.1016/0749-596X(86)90024-0
Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005
Belletti, A., & Rizzi, L. (1988). Psych-verbs and q -Theory. *Natural Language and Linguistic Theory*, *6*(3), 291–352. https://doi.org/10.1007/BF00133902
Bosch, P., & Umbach, C. (2007). Reference determination for demonstrative pronouns. *ZAS Papers in Linguistics*, *48*, 39–51. https://doi.org/10.21248/zaspil.48.2007.353
Bott, O., & Solstad, T. (2014). From verbs to discourse: A novel account of implicit causality. In C. Fabricius-Hansen, B. Hemforth, & B. Mertins (Eds.), *Psycholinguistic approaches to meaning and understanding across languages* (pp. 213–251). Springer.
Bresnan, J., & Kanerva, J. M. (1989). Locative inversion in Chichewa: A case study of factorization in grammar. *Linguistic Inquiry*, *20*(1), 1–50.
Brilmayer, I., & Schumacher, P. B. (2021). Referential chains reveal predictive processes and form-to-function mapping: An electroencephalographic study using naturalistic story stimuli. *Frontiers in Psychology*, *12*, 623648. https://doi.org/10.3389/fpsyg.2021.623648
Burnsky, J., Keshev, M., Asatryan, M., Hlachova, B., Johnson, K., & Dillon, B. (2022, March 24–26). *Look away! An object is coming* [Poster presentation]. The 35th Annual Conference on Human Sentence Processing, UC Santa Cruz, USA.
Caramazza, A., Grober, E., Garvey, C., & Yates, J. (1977). Comprehension of anaphoric pronouns. *Journal of Verbal Learning and Verbal Behavior*, *16*(5), 601–609. https://doi.org/10.1016/S0022-5371(77)80022-4
Chambers, C. G., & Smyth, R. (1998). Structural parallelism and discourse coherence: A test of centering theory. *Journal of Memory and Language*, *39*(4), 593–608. https://doi.org/10.1006/jmla.1998.2575
Colonna, S., Schimke, S., & Hemforth, B. (2012). Information structure effects on anaphora resolution in German and French: A cross-linguistic study of pronoun resolution. *Linguistics*, *50*(5), 991–1013. https://doi.org/10.1515/ling-2012-0031
Colonna, S., Schimke, S., & Hemforth, B. (2014). Information structure and pronoun resolution in German and French: Evidence from the visual-world paradigm. In B. Hemforth, B. Schmiedtová, & C. Fabricius-Hansen (Eds.), *Psycholinguistic approaches to meaning and understanding across languages* (pp. 175–195). Springer.
Cowles, H. W. (2007). Influence of 'aboutness' on pronominal coreference. *ZAS Papers in Linguistics*, *48*, 23–38. http://dx.doi.org/10.21248/zaspil.48.2007.352
Crawley, A., Stevenson, J., & Kleinman, D. (1990). The use of Heuristic Strategies in the interpretation of pronouns. *Journal of Psycholinguistic Research*, *19*(4), 245–264. https://doi.org/10.1007/BF01077259
Crawley, R. A., & Stevenson, R. J. (1990). Reference in single sentences and in texts. *Journal of Psycholinguistic Research*, *19*(3), 191–210. https://doi.org/10.1007/BF01077416
Creemers, A., & Meyer, A. S. (2022). The processing of ambiguous pronominal reference is sensitive to depth of processing. *Glossa Psycholinguistics*, *1*(1). https://doi.org/10.5070/G601166
Dixon, R. M. (1979). Ergativity. *Language*, *55*(1), 59–138. https://doi.org/10.2307/412519
Do, M. L., & Kaiser, E. (2022). Sentence formulation is easier when thematic and syntactic prominence align: Evidence from psych verbs. *Language, Cognition and Neuroscience*, *37*(5), 648–670. https://doi.org/10.1080/23273798.2021.2008458
Ellert, M. (2013). Resolving ambiguous pronouns in a second language: A visual-world eye-tracking study with Dutch learners of German. *International Review of Applied Linguistics in Language Teaching*, *51*(2), 171–197. https://doi.org/10.1515/iral-2013-0008
Fedriani, C. (2014). *Experiential constructions in Latin*. Brill.
Ferreira, F. (1994). Choice of passive voice is affected by verb type and animacy. *Journal of Memory and Language*, *33*(6), 715–736. https://doi.org/10.1006/jmla.1994.1034
Ferstl, E. C., Garnham, A., & Manouilidou, C. (2011). Implicit causality bias in English: A corpus of 300 verbs. *Behavior Research Methods*, *43*(1), 124–135. https://doi.org/10.3758/s13428-010-0023-2
Foley, W. (2005). Semantic parameters and the unaccusative split in the Austronesian language family. *Studies in Language*, *29*(2), 385–430.

Foley, W., & Van Valin, R. (1984). *Functional syntax and universal grammar*. Cambridge University Press.

Fuchs, M., & Schumacher, P. B. (2020). Referential shift potential of demonstrative pronouns–evidence from text continuation. In Å Næss, A. Margetts, & Y. Treis (Eds.), *Demonstratives in discourse* (pp. 185–213). Language Science Press.

Gernsbacher, M. A., & Hargreaves, D. J. (1988). Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language*, *27*(6), 699–717. https://doi.org/10.1016/0749-596X(88)90016-2

Gernsbacher, M. A., Hargreaves, D. J., & Beeman, M. (1989). Building and accessing clausal representations: The advantage of first mention versus the advantage of clause recency. *Journal of Memory and Language*, *28*(6), 735–755. https://doi.org/10.1016/0749-596X(89)90006-5

Giorgi, A. (1984). Towards a theory of long distance anaphors: A GB approach. *The Linguistic Review*, *3*(4), 307–361. https://doi.org/10.1515/tlir.1984.3.4.307

Givón, T. (1983). *Topic continuity in discourse: Quantified cross-language studies*. John Benjamins.

Gordon, P. C., Grosz, B. J., & Gilliom, L. A. (1993). Pronouns, names, and the centering of attention in discourse. *Cognitive Science*, *17*(3), 311–347. https://doi.org/10.1207/s15516709cog1703_1

Grimshaw, J. (1990). *Argument structure*. The MIT Press.

Grober, E. H., Beardsley, W., & Caramazza, A. (1978). Parallel function strategy in pronoun assignment. *Cognition*, *6*(2), 117–133. https://doi.org/10.1016/0010-0277(78)90018-5

Grosz, B. J., Joshi, A. K., & Weinstein, S. (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, *21*(2), 203–225.

Gundel, J. K., Hedberg, N., & Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, *69*(2), 274–307. https://doi.org/10.2307/416535

Hartshorne, J. K., O'Donnell, T. J., & Tenenbaum, J. B. (2015). The causes and consequences explicit in verbs. *Language, Cognition and Neuroscience*, *30*(6), 716–734. https://doi.org/10.1080/23273798.2015.1008524

Hartshorne, J. K., & Snedeker, J. (2013). Verb argument structure predicts implicit causality: The advantages of finer-grained semantics. *Language and Cognitive Processes*, *28*(10), 1474–1508. https://doi.org/10.1080/01690965.2012.689305

Hobbs, J. (1979). Coherence and coreference*. *Cognitive Science*, *3*(1), 67–90. https://doi.org/10.1207/s15516709cog0301_4

Hopper, P. J., & Thompson, S. A. (1980). Transitivity in grammar and discourse. *Language*, *56*(2), 251–299. https://doi.org/10.1353/lan.1980.0017

Johnson, E., & Arnold, J. E. (2021). Individual differences in print exposure predict use of implicit causality in pronoun comprehension and referential prediction. *Frontiers in Psychology*, *12*, 672109. https://doi.org/10.3389/fpsyg.2021.672109

Kaiser, E. (2009). Investigating effects of structural and information-structural factors on pronoun resolution. In M. Zimmermann, & C. Féry (Eds.), *Information structure from different perspectives* (pp. 332–353). Oxford University Press.

Kaiser, E. (2011a). Focusing on pronouns: Consequences of subjecthood, pronominalisation, and contrastive focus. *Language and Cognitive Processes*, *26*(10), 1625–1666. https://doi.org/10.1080/01690965.2010.523082

Kaiser, E. (2011b). On the relation between coherence relations and anaphoric demonstratives in German. In I. Reich, E. Horch, & D. Pauly (Eds.), *Proceedings of Sinn & Bedeutung* (Vol. 15, pp. 337–351). Saarland University Press.

Kaiser, E. (2019). Linguistic consequences of event segmentation in visual narratives: Implications for prominence. *Language, Cognition and Neuroscience*, *35*(3), 402–408. https://doi.org/10.1080/23273798.2019.1667000

Kaiser, E., & Trueswell, J. C. (2008). Interpreting pronouns and demonstratives in Finnish: Evidence for a form-specific approach to reference resolution. *Language and Cognitive Processes*, *23*(5), 709–748. https://doi.org/10.1080/01690960701771220

Kehler, A. (2002). *Coherence, reference, and the theory of grammar*. CSLI Publications.

Kehler, A., Kertz, L., Rohde, H., & Elman, J. (2008). Coherence and coreference revisited. *Journal of Semantics*, *25*(1), 1–44. https://doi.org/10.1093/jos/ffm018

Kehler, A., & Rohde, H. (2013). A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics*, *39*(1-2), 1–37. https://doi.org/10.1515/tl-2013-0001

Kertz, L., Kehler, A., & Elman, J. (2006). Grammatical and coherence-based factors in pronoun interpretation. *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, *28*, 1605–1610.

Koornneef, A. W., & Sanders, T. J. (2013). Establishing coherence relations in discourse: The influence of implicit causality and connectives on pronoun resolution. *Language and Cognitive Processes*, *28*(8), 1169–1206. https://doi.org/10.1080/01690965.2012.699076

Koornneef, A. W., & Van Berkum, J. J. A. (2006). On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye tracking. *Journal of Memory and Language*, *54*, 445–465. https://doi.org/10.1016/j.jml.2005.12.003

McKoon, G., Greene, S. B., & Ratcliff, R. (1993). Discourse models, pronoun resolution, and the implicit causality of verbs. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(5), 1040–1052. https://doi.org/10.1037/0278-7393.19.5.1040

Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. MIT press.

Patterson, C., Schumacher, P. B., Nicenboim, B., Hagen, J., & Kehler, A. (2022). A Bayesian approach to German personal and demonstrative pronouns. *Frontiers in Psychology*, *12*, 672927. https://doi.org/10.3389/fpsyg.2021.672927

Qualtrics. (2019). *Qualtrics* [Computer software]. https://www.qualtrics.com

Qualtrics. (2020). *Qualtrics* [Computer software]. https://www.qualtrics.com

Rappaport Hovav, M., & Levin, B. (2015). The syntax-semantics interface: Semantic roles and syntactic arguments. In S. Lappin, & C. Fox (Eds.), *The handbook of contemporary semantic theory* (pp. 593–634). Blackwell.

R Core Team. (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. URL https://www.R-project.org/

Rohde, H. (2008). *Coherence-driven effects in sentence and discourse processing* [Unpublished doctoral dissertation]. University of California, San Diego.

Rohde, H., & Kehler, A. (2008, March 13–15). *The bidirectional influence between coherence establishment and pronoun interpretation* [Poster presentation]. The 21st Annual CUNY

Conference on Human Sentence Processing, Chapel Hill, NC, USA.

Rohde, H., Kehler, A., & Elman, J. L. (2006). Event structure and discourse coherence biases in pronoun interpretation. *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*, *28*, 697–702.

Rudolph, U., & Försterling, F. (1997). The psychological causality implicit in verbs: A review. *Psychological Bulletin*, *121*(2), 192–218. https://doi.org/10.1037/0033-2909.121.2.192

Runner, J. T., & Ibarra, A. (2016). Information structure effects on null and overt subject comprehension in Spanish. In A. Holler & K. Suckow (Eds.), *Empirical perspectives on anaphora resolution: Information structural evidence in the race for salience* (pp. 1–69). de Gruyter. http://dx.doi.org/10.1515/9783110464108-006

Schumacher, P. B., Dangl, M., Uzun, E., Holler, A., & Suckow, K. (2016). Thematic role as prominence cue during pronoun resolution in German. In A. Holler, C. Göb, & K. Suckow (Eds.), *Empirical perspectives on anaphora resolution: Information structural evidence in the race for salience* (pp. 121–147). de Gruyter.

Schumacher, P. B., Roberts, L., & Järvikivi, J. (2017). Agentivity drives real-time pronoun resolution: Evidence from German er and der. *Lingua. International Review of General Linguistics. Revue internationale De Linguistique Generale*, *185*, 25–41. https://doi.org/10.1016/j.lingua.2016.07.004

Smyth, R. (1994). Grammatical determinants of ambiguous pronoun resolution. *Journal of Psycholinguistic Research*, *23*(3), 197–229. https://doi.org/10.1007/BF02139085

Song, J., & Kaiser, E. (2020a). Effects of discourse factors on the interpretation of Korean null pronouns in subject and object position. *Proceedings of the Linguistic Society of America*, *5*(*1*), 655. http://dx.doi.org/10.3765/plsa.v5i1.4751

Song, J., & Kaiser, E. (2020b). Forward-looking effects in subject pronoun interpretation: What comes next matters. *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society*, *42*, 702–708.

Speas, M. J. (1990). *Phrase structure in natural language (Studies in Natural Language and Linguistic Theory)* (Vol. 21). Kluwer.

Stevenson, R., Nelson, A., & Stenning, K. (1995). The role of parallelism in strategies of pronoun comprehension. *Language and Speech*, *38*(4), 393–418. https://doi.org/10.1177/002383099503800404

Stevenson, R. J., Crawley, R. A., & Kleinman, D. (1994). Thematic roles, focus and the representation of events. *Language and Cognitive Processes*, *9*(4), 519–548. https://doi.org/10.1080/01690969408402130

Stewart, A. J., Holler, J., & Kidd, E. (2007). Shallow processing of ambiguous pronouns: Evidence for delay. *Quarterly Journal of Experimental Psychology*, *60*(12), 1680–1696. https://doi.org/10.1080/17470210601160807

Ueno, M., & Kehler, A. (2016). Grammatical and pragmatic factors in the interpretation of Japanese null and overt pronouns. *Linguistics*, *54*(6), 1165–1221. https://doi.org/10.1515/ling-2016-0027

Van Valin, R. D. (2001). *An introduction to syntax*. Cambridge University Press.

Van Valin, R. D., Jr., & LaPolla, R. (1997). *Syntax: Structure, meaning, and function*. Cambridge University Press.

Verhoeven, E. (2009). Subjects, agents, experiencers, and animates in competition: Modern Greek argument order. *Linguistische Berichte*, *2009*(219), 355–376.

Verhoeven, E. (2014). Thematic prominence and animacy asymmetries. Evidence from a cross-linguistic production study. *Lingua. International Review of General Linguistics. Revue internationale De Linguistique Generale*, *143*, 129–161. https://doi.org/10.1016/j.lingua.2014.02.002

Walker, M. A., Joshi, A. K., & Prince, E. F. (1998). Centering in naturally occurring discourse: An overview. In M. A. Walker, A. K. Joshi, & E. F. Prince (Eds.), *Centering theory in discourse* (pp. 1–28). Clarendon Press.

Winograd, T. (1972). Understanding natural language. *Cognitive Psychology*, *3*(1), 1–191. https://doi.org/10.1016/0010-0285(72)90002-3

Wolf, F., Gibson, E., & Desmet, T. (2004). Discourse coherence and pronoun resolution. *Language and Cognitive Processes*, *19*(6), 665–675. https://doi.org/10.1080/01690960444000034