

Beyond Principle C: (Not)-at-issueness and plausibility influence acceptability of coconstrual

Vera Gor and Kristen Syrett
Rutgers University

1 Introduction

Binding Theory, as originally proposed in Chomsky (1981), imposes restrictions on the distribution of nominal expressions that are intended to pick out the same referent. Principle C of the Binding Theory, in particular, requires that a name, or an R-expression, must be free. That is, it must not have a co-indexed c-commanding antecedent. As a result, Principle C accounts for a range of ungrammatical data, including sentences like those in (1) - (3), where the name *Susan* is located in the c-commanding domain of a pronoun or another name referring to the same individual.

- (1) ***She_i** likes **Susan_i**.
- (2) ***Susan_i** likes **Susan_i**.
- (3) ***She_i** always goes to Millennium Park when **Susan_i** comes to Chicago.

At the same time, speakers occasionally allow for a pronoun and a name in its c-commanding domain to refer to the same individual, despite the constraint of Principle C. Indeed, over the years, researchers investigating the syntax and semantics of binding relations have collected and discussed a significant number of such apparent counterexamples to Principle C, as illustrated in (4) - (7).

- (4) **He_i** did what **John_i** always does. (Bolinger 1979: 292 (30), cited in Sag 2000: (17a))
- (5) **He_i** became known as **Napoleon_i**. (Levinson 2000: 302 (40j))
- (6) **He_i** put on **John_i**'s coat; but only John would do that; so **he_i** is **John_i**. (Higginbotham 1985: 570 (63), adapted and further discussed in Safir 2004: 28 (7))
- (7) Everyone has finally realized that Oscar is incompetent. Even **he_i** has finally realized that **Oscar_i** is incompetent. (Evans 1980: 357 (52), also cited in Grodzinsky & Reinhart 1993: 78 (19c))

Examples like these and others seem to indicate that the traditional formulation of Binding Principle C falls short of providing an accurate account of just where and when speakers allow for a pronoun and a name in its c-commanding domain to refer to one and the same individual. Thus, Principle C alone is insufficient to determine which coconstrual relations will be licensed by speakers and which will not be.¹

¹Following Safir (2004), we use the term *coconstrual* to broadly refer to a full identity overlap relation between two nominals, without necessarily specifying whether this identity overlap is a result of a dependent identity relation or a coreference relation.

Addressing such problematic data, a number of researchers in the last few decades have looked outside of the domain of syntax to pragmatics. However, each proposal presents challenges of its own. For example, Heim (1982) appealed to the notion of guises, proposing that the same individual may be represented by more than one index if the pragmatics of the situation is such that it supplies distinct perspectives on this individual. Heim further argued that nominals pick out mental rather than physical representations and that a minimal unit of reference in fact corresponds to one of potentially several guises or mental representation of the same individual (Heim 1982). While this proposal could potentially provide an account of data such as in (4), it does not offer a precise formulation of what specific conditions need to be observed so that the *pragmatics of the situation* would allow for viewing an individual as represented by two distinct guises.

Reinhart (1983) argued that the constraint on the distribution of intrasentential coreference is not syntactic, i.e., it is not governed by an independent condition of the grammar, but follows pragmatically from the bound anaphora conditions. More specifically, Reinhart (1983) views condition on coreference as an instance of Gricean generalized implicatures: in a rational linguistic exchange, one would expect that if a speaker has the means to express a certain idea clearly and directly, he/she would not choose, arbitrarily, a less clear way to express it. This proposal provides an account of such cases as (5)-(7), where the same idea could not potentially be expressed by switching the positions of pronoun and name in the sentence. Still it falls short of offering an explanation of data such as in (4).

Arguing for a division of labor between syntax and semantics/pragmatics, Chien & Wexler (1990) proposed that binding principles B and C always impose conindexing. To prevent the possibility of accidental coreference where two distinct indices accidentally pick out the same referent, Chien & Wexler (1990) proposed pragmatic Principle P, which requires conindexed NPs to be non-coreferential, unless the context explicitly forces coreference. This proposal restricts coconstrual and allows for exceptions based on context, and may provide an account for identity sentences, such as (5), which typically emerge in contexts that introduce an identity relation that is not obvious to the addressee. However, a potential problem is that it does not specify the properties that a context needs to have in order to impose or allow coreference on the two conindexed nominals.

Finally, Safir (2004) also appealed to non-structural factors to account for apparent counterexamples to Principle C. He reformulated Principle C as *pragmatic obviation*, arguing that when a name is c-commanded by a pronoun, this structural relation between the two nominals creates an *expectation* of non-coconstrual. This expectation can be further adjusted by pragmatic context. Formulated this way, *pragmatic obviation* can account for cases such as (5) - (6), and *instantiation contexts*, such as (7) – cases where individuals are singled out as instantiations of properties under discussion. Still, Safir (2004), just as Chien & Wexler (1990), does not offer a discussion of which specific pragmatic properties of the context can lead to altering one's expectations of non-coconstrual.

The situation we face is this. Chomsky's Principle C simply does not allow for coconstrual between a pronoun and a name in its c-commanding domain, and therefore drastically undergenerates cases when such a relation would be deemed permissible by native speakers. At the same time, accounts that appeal to various non-

structural, pragmatic factors to account for counterexamples where speakers find coconstrual acceptable despite a restriction imposed by Principle C lack a level of specificity that would allow us to systematically predict whether future instances of attempted coconstrual in the face of a Principle C effects should be licensed or not, and therefore overgenerate acceptability inferences. So the main question motivating the current paper is, *What are the specific pragmatic or contextual factors that influence speakers' judgments of acceptable, but structurally illicit, coconstrual?*

In our previous and ongoing experimental research on the factors licensing what we will term "structurally illicit, pragmatically viable coconstrual," we have focused on the environment of backwards anaphora in a variety of instantiations to demonstrate the robustness of the phenomenon and the subtle manipulations that give rise to such acceptability patterns. In these studies, we have uncovered three main findings thus far using a forced-choice paradigm in which participants are asked to choose between an intra- and an extra-sentential referent for a pronoun (as they are in the experiment reported below). In the data that follow, the percentage of time participants chose an intrasentential referent (therefore allowing coconstrual between the pronoun and the name it c-commands) is indicated to the right.

The first factor that has been found to have an effect on acceptability is the structural position of the c-commanding pronoun, as shown in (8). When the pronoun is in subject position, as in (8a), coconstrual is disallowed. By contrast, when the pronoun is in object position, as in (8b), coconstrual is allowed. However, there is a caveat, which is illustrated in our second finding. Participants' judgments of acceptability are systematically influenced by the conceptual plausibility of a particular interpretation, which we have argued taps into a stored schema in memory (Gor & Syrett 2018). For example, in (9a), it is unlikely that Mark is offering Emily her own book to read, and hence coconstrual is not allowed. However, it *is* likely that a waiter might offer Emily (who is perhaps a regular patron at the restaurant) her (own) favorite entrée, and so coconstrual is allowed in (9b) – at a rate *much* higher than structural principles alone would predict. Likewise, in (8b) it is likely for Mr. Tomkins to show Emily the new desk at which Emily will be working. A significant amount of experimental research has shown that manipulating plausibility influences parsing and has an effect on resolving ambiguities (Boland *et al.* 1995, Clifton 1993, Pickering & Traxler 1998, Tanenhaus *et al.* 1989, Traxler & Pickering 1996, Trueswell *et al.* 1993, a.m.o.). Finally, by keeping the structural position of the pronoun constant (i.e., maintaining its c-command position relative to the name and holding it constant as an object in the sentence) while manipulating the position of the pronoun relative to that of a comparative morpheme in a comparative construction, we have found that coconstrual is more acceptable in subject comparatives (10b) than object comparatives (10a) (Gor 2017, Gor & Syrett 2015).

(8) *Structural position of the pronoun:*

- a. **She**_{*i*} showed Max **Emily**_{*i*}'s diary. (0%)
- b. Mr. Tomkins showed **her**_{*i*} **Emily**_{*i*}'s new desk. (32%)

(9) *Plausibility of the coconstrual relation:*

- a. Mark offered **her**_{*i*} **Emily**_{*i*}'s book to read. (3%)
- b. The waiter offered **her**_{*i*} **Emily**_{*i*}'s favorite entrée. (38%)

(10) *Order of operations:*

- a. Friends sent **her_i** a cheaper gift than they sent **Marianna_i** last year. (10%)
- b. More relatives wanted **her_i** to study abroad than at **Amelia_i's** community college. (64%)

We have argued that in the case of both the difference in pronominal structural position and with comparatives, there may be a common factor leading to increased acceptability: the order of operations in the incremental processing of the sentence. Specifically, when the processor encounters a pronoun in sentence-initial position, it immediately activates Principle C, barring coconstrual between the pronoun and the name that follows. By contrast, when the pronoun is in object position, the processor has already been engaged in processing the initial part of the sentence, leaving the door open for other information relevant to determining coconstrual relations to play a role.² In the same vein, when the processor initially encounters a comparative (*more*), this cues the processor to launch a comparison of alternatives, overshadowing the binding constraints that would otherwise disfavor coconstrual. Note that in all cases where coconstrual is deemed acceptable, Principle C is not completely out of the picture: percentages appear to max out at a certain level. However, this variability in percentages under carefully controlled conditions demonstrates the importance of information structure and pragmatics in interacting with Principle C in determining coconstrual. These findings are in line with independent evidence from the processing of presuppositions signaled by hard triggers such as *too* and *auch* in German that the representation of the context is rapidly updated as soon as possible, even upon encountering a noun phrase, or (here) a comparative marker. (See also Schwarz (2007) and Schwarz & Tiemann (2017)).

Following in the footsteps of the previous research, we continue to probe the influence of plausibility, information status, and processing on coconstrual, extending our focus to the at-issue, projective status of the proposition containing the pronoun and name. Content that is *at issue* (AI) addresses the Question Under Discussion (Roberts 1996), and *does not* project past operators (Potts 2005, Simons *et al.* 2010, Tonhauser 2012). By contrast, content that is *not at issue* (NAI) does not address the QUD, and projects. As a result of its NAI status, it resists direct rejection.

A paradigmatic example of NAI content is *Appositive Relative Clauses* (also known as Non-Restrictive Relative Clauses), as in the bold portion of (11) (Ander-Bois *et al.* 2013, Harris & Potts 2009, Potts 2005, Schlenker 2010, Syrett & Koev 2014, Tonhauser 2012, a.o.).

- (11) My friend Sophie, **who is a classical violinist**, performed a piece by Mozart.
(from Syrett & Koev 2014: 1 (1))
- a. *At-Issue content*: My friend Sophie performed a piece by Mozart.
 - b. *Not-At-Issue content*: My friend Sophie is a classical violinist.

²Previous psycholinguistic research has provided compelling evidence that the parser categorically ignores antecedents in structurally illicit positions (Kazanina *et al.* 2007). However, the stimuli in this study placed the pronoun in sentence-initial position, and we therefore put forth that the particular stimuli in this research led to this claim. When a broader range of stimuli are considered, the picture of Principle C becomes more nuanced.

Appositives are right-adjoined to their nominal anchor and express a proposition (which does make truth conditional contributions comparable to that of coordinated conjuncts (Syrett & Koev 2014)), but this proposition presents content that is treated differently from the main clause in that it is not at issue. Note that a rejection of *No, she isn't!* seems odd, while *No, she didn't!* seems acceptable.³

Dillon *et al.* (2014) and Dillon *et al.* (2017) have shown experimentally in both online and offline studies with appositives that NAI content of the utterance is processed differently than AI content, impeding the processing of complex or lengthy syntactic material (in particular, the tracking of *wh*- filler gap dependencies) less, but treating violations of grammatical number agreement (as shown in (12)) the same. This contrast is critical to note: while syntactic complexity is processed differently in NAI versus AI content, the processor does not distinguish between hard-wired, categorical grammatical violations in the two sources: if it's bad, it's bad, NAI-ness does not salvage it.

(12) Sample stimuli from (Dillon *et al.* (2014)), Experiment 2:

- a. *RRC, agreement mismatch*:
That economist who was at the conference *Paul attend* every year in Rome made fun of the banker.
- b. *ARC, agreement mismatch*:
That economist, the one who was at the conference *Paul attend* every year in Rome, made fun of the banker

Appositives are, of course, not the only type of construction representing NAI content. In fact, all presupposition triggers signal NAI content that projects. For example, temporal clauses are known to trigger presuppositions. They also pass two key tests for presenting NAI content: the now-standard diagnostics of attempting to target the material with a direct rejection (Tonhauser 2012, Syrett & Koev 2014), as in (13), and attempting to access the main assertion with a *Why?* question, as in (14). The temporal clause fails to provide the basis for a felicitous direct rejection or as a reason in response to the question.

(13) After Richard wrapped the presents, Amelia put on the bows.

- a. #No, he didn't!
- b. No, she didn't!

(14) After Sarah took some Advil, she fixed herself a sandwich.

Why?

- a. #Because she had a headache.
- b. Because she was hungry.

We have conducted an independent study in the lab involving target sentences just like these without any pronominal anaphora, to confirm the robust NAI status of temporal clauses. Participants overwhelmingly preferred to target the matrix, not

³We leave aside the difference between medial and final appositives here to focus on the basic NAI/AI distinction. See AnderBois *et al.* (2013), Gobel (2018), Hunter & Asher (2016), and Syrett & Koev (2014) for discussion.

the temporal clause, with a direct rejection: 97.8% chose to reject the matrix, not the appositive clause in control sentences, vs. 96.2% rejecting the content of the matrix, not the temporal clause in the target items. Further 95.9% chose not to target the temporal clause with a *Why?* question. This point of commonality between appositives and the presupposition-triggering environment of a temporal clause presents us with an opportunity to investigate the influence of NAI/AI status of the potential coconstrual relation while manipulating a known entity (plausibility).

We investigate appositives in a parallel line of research. Here, we turn to sentence-initial temporal clauses for three main reasons. First, they allow us to place the target construction hosting the potential coconstrual relationship in sentence-initial position while holding the pronoun constant in object position, thereby satisfying one precondition on paving the way for coconstrual. Second, the sentence-initial presence of the adverb heading the clause ("after") should signal to the processor that it should begin updating the context set, and thus drawing attention away from the binding conditions, and should also trigger the presupposition and consequently the processing of NAI content. Finally, we are easily able to insert our previously-tested stimuli into a sentence-initial adverbial clause in order to contrast the percentage of coconstrual we previously obtained with the percentage we get when the exact same relation is inserted into a NAI temporal clause, as illustrated in (15), where the original stimulus is in (15a) (repeated from (9b) and the same target embedded into a temporal clause is shown in (15b).

(15) *Difference in location of possible coconstrual based on (N)AI status:*

- a. The waiter offered **her_i Emily_i's** favorite entrée.
- b. After offering **her_i Emily_i's** favorite entrée, the waiter brought a pitcher of water.

The relevant question we pursue here is as follows. Does the NAI status of the temporal clause affect participants' judgments of syntactically disfavored backwards anaphora contained in this clause, elevating acceptability of coconstrual not only above levels predicted by Principle C but also above levels where the coconstrual relation is contained in AI content? We can formulate two competing hypotheses.

(H₀) *Default Hypothesis: no NAI/AI difference* - If, following Chomsky (1981)'s formulation of the binding constraints, Principle C is to be interpreted as a hard and fast principle of the grammar, judged in a vein similar to morphosyntactic agreement, then treatment of structurally illicit coconstruals will be on par across AI and NAI content, and we would observe the same rate of rejection in both types of environments.

(H_A) *Alternative Hypothesis: NAI/AI difference* - If, as the evidence reviewed above suggests, Principle C is not a deterministic enforcer of binding constraints, but rather a factor that interacts with others within and outside of the grammar, then participants' treatment of coconstrual will depend on the (N)AI status of the content in which the pronoun-name sequence is embedded, and given differential processing of NAI/AI content, coconstrual will be more likely in NAI content, *modulo* degree of plausibility.

In this paper, we report the findings of an experiment designed to test these hypotheses. The results show that participants accept coconstrual in backwards

anaphora in which Principle C is implicated when plausibility of coconstrual is high (as predicted by our previous work), and even more so when such structurally illicit coconstruals are introduced as part of NAI content of the utterance (i.e., in a sentence-initial temporal clause). The results therefore provide further reason to reconsider the status of Principle C in the grammar and the ways in which the processor implements constraints on coconstrual relations.

2 Experiment

We present an offline, forced-choice task (a paradigm we have successfully used earlier) to assess whether the (not)-at-issue status of the content in which a structurally illicit coconstrual relation appears influences participants' assessment of the coconstrual relation between a pronoun and a name it c-commands. We further assess the additional role of plausibility as an independent and interacting factor.

2.1 Participants

82 Rutgers University undergraduates, all native speakers of English (as determined by a demographic questionnaire), participated. All subjects gave informed consent and received course credit for their participation. Data from 35 additional participants were excluded, because they were not native speakers of English (n=21) or they failed more than 10% of control sentences (3 and more out of 24) (n=14).

2.1.1 Design

All target sentences included a pronoun *she* c-commanding a female name. The name was embedded in a possessive DP (e.g., *Emily's grandfather* or *Pamela's notes*) to allow for the use of predicates that require two distinct individuals/entities as their arguments in the target sentences. In all target stimuli the pronoun therefore not only linearly preceded the name, but also structurally dominated it, thus yielding a Principle C effect. The experiment had a 2x2 design with two factors manipulated (each with two levels):

(a) *AI status* of the clause (proposition) hosting the pronoun-name sequence (NAI sentence-initial adjunct temporal clause vs. AI matrix clause), and

(b) the *plausibility of coconstrual* between the pronoun and the name (low vs. high). Plausibility was determined by the rankings obtained in an independent norming study reported in Gor & Syrett (2018). See that paper for more details of the norming study.

An example of a set of target sentences is presented in (16), with the pronoun and name in bold. We predicted that the target sentences such as in (16d) would elicit the highest levels of coconstrual.

- (16) Sample target items with factors of *AI status* and *plausibility of coconstrual* crossed
- a. *AI, low plausibility*:
Mr. Adams allowed **her_i** to borrow **Emily_i's** notes for the exam.
 - b. *AI, high plausibility*:
The doctors allowed **her_i** to visit **Emily_i's** grandfather in the ICU.

- c. *NAI, low plausibility:*
After allowing **her_i** to borrow **Emily_i's** notes for the exam, Mr. Adams phoned the library about the new textbook.
- d. *NAI, high plausibility:*
After allowing **her_i** to visit **Emily_i's** grandfather in the ICU, the doctors discussed the case with the radiologist.

Following the design from Gor & Syrett (2018), in each target sentence, the clause that incorporated a structurally illicit coconstrual featured either an exceptional case marking (ECM) or a double object (DO) predicate. This choice of construction was motivated by the fact that both types of predicates allow for an argument position lower than the matrix subject to c-command the remaining linguistic material in the clause, i.e., an embedded subject position in an ECM construction, and an indirect object position in a DO construction. Thus it was made possible for the pronoun to dominate the name from a position in the clause lower than the matrix subject position. In the AI versions, this matrix subject was present. In the NAI versions, the subject was absent. However, crucially, the structural relation between the pronoun (*her*) and the name (*Emily*) was kept identical across both conditions. Moreover, the distribution of thematic roles was also preserved: in both conditions, the matrix subject DP (e.g., *the doctors*) was the agent (e.g., of the predicate *allow*); and the pronoun *her* was the benefactor of the same verb.

Altogether there were 12 pairs of NAI/AI target items with high plausibility of coconstrual and 10 pairs of NAI/AI target items with low plausibility of coconstrual, yielding a total of 44 target sentences. These sentences were distributed across four lists so that each participant saw only one sentence from each set and a total of 11 target sentences. There were also 24 control items, all of which involved forward anaphora with no Principle C effects, as shown in (17). Plausibility was also manipulated among control items, as shown in the contrast between (17a) and (17b), to elicit a baseline for the influence of plausibility in the absence of any Principle C effects. These 11 targets and 24 controls were pseudorandomized with 45 filler sentences, for a total of 80 items per participant.

- (17) Sample control items with forward anaphora
 - a. *High plausibility:*
Emily's friends were planning a surprise birthday party for **her**.
 - b. *Low plausibility:*
Pamela was invited to **her** exhibition opening.

2.2 Procedure

The experiment took place in a quiet laboratory setting, with participants run one or two at a time at individual response stations. Stimuli were presented on a 21-inch iMac screen using Cedrus SuperLab 5 stimulus presentation software (SuperLab Pro 2015). The experiment began with a brief training session, to acclimate participants to the instructions in the task and the use of the response pad, and to increase the salience of two potential female antecedents who were visually depicted on the screen and who were referenced in the linguistic stimuli.

During the experiment proper, each trial had the same structure. A target sentence was displayed, accompanied by the images of two female characters labeled as Emily and Pamela, as in Fig. 1. The L/R positions of the two girls remained constant throughout the task. Appearance of the name was balanced across target and control sentences, to vary responses and expectations.




Figure 1: Sample stimulus item.

Participants were instructed to read each sentence, and choose which one of the two female characters the pronoun in the sentence referred to. In this way, they were implicitly asked to choose between an intra-sentential or an extra-sentential referent for the pronoun (where the former would indicate a licit coconstrual relation, regardless of structural constraints). Participants indicated their answers by pressing a key labeled with either letter E or letter P on the response pad. The experimental session lasted approximately 15-20 minutes.

2.3 Results

Results for control items and target items are presented in Fig. 2 below. We begin by discussing the results of the controls, which featured forward anaphora and therefore no Principle C effects. We thus predicted that the choice of referent would primarily be guided by plausibility, and indeed, our findings bore out this prediction. 98.8% of participants chose an intra-sentential antecedent for a pronoun in sentences when plausibility of coconstrual was high, while the choice of intra-sentential antecedent for low plausibility items was only 10.7%. This pattern is in line with the results reported in Gor & Syrett (2018) and therefore provides further evidence for the robust influence of plausibility in resolving pronominal ambiguity, regardless of structural constraints.

We now turn to the target items. As expected from Gor & Syrett (2018), backwards anaphora sentences with low-ranked plausibility of coconstrual yielded a correspondingly low percentage of intra-sentential referent selection (2.09% in AI condition, and 7.93% in NAI condition). By contrast, those sentences with a high level of coconstrual plausibility yielded percentages that were higher than predicted by the c-command relations between the pronoun and the name. The effect was most

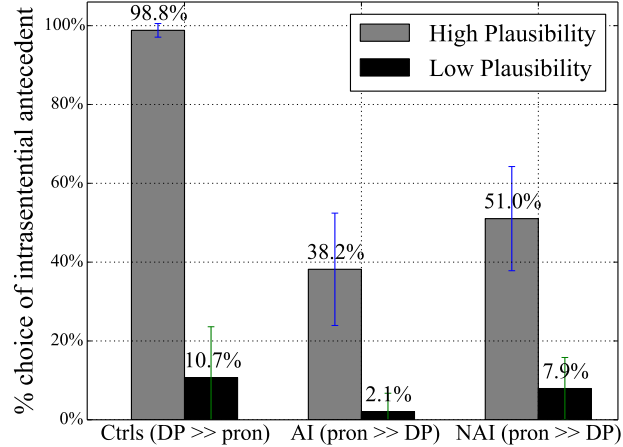


Figure 2: Percentage choice of selection of intra-sentential antecedent for the pronoun across conditions.

pronounced in cases where the pronoun-name sequence was introduced as NAI content in a sentence-initial temporal clause (51.04% in NAI vs. 38.18% AI).

We performed a binomial logistic regression model with subjects and items as random intercepts. The analysis revealed significant effects of both plausibility ($\beta = 3.6728$, $SE = 0.4920$, $p < 0.001$) and AI/NAI status of the proposition ($\beta = -1.5790$, $SE = 0.6892$, $p < 0.05$). There was no significant interaction between the two factors ($\beta = 0.8044$, $SE = 0.7771$, $p = 0.301$), suggesting that NAI status of the proposition had the same impact of increasing coconstrual across the board.

To probe the results further, we turned to an analysis of participants' responses to individual target items. As shown in Fig. 3, for 19 out of 22 target items (86.4%), the percentage of participants choosing the intra-sentential antecedent in the NAI condition exceeded the respective percentage in the AI condition. This pattern was observed for 100% of target items with low plausibility of coconstrual and 75% of target items with high plausibility of coconstrual (items in the shaded area in Fig. 3). The clustering of the low plausibility items in the lower left corner in contrast to the variability observed with the high plausibility items, with most gravitating towards high levels of coconstrual, regardless of (N)AI status, reinforces that role of plausibility in paving the way for coconstrual.

We then turned to analyzing individual participant responses, homing in on the high plausibility condition, where coconstrual was a viable option. In Fig. 4 and Fig. 5, we present histograms of individual responses to the AI items and NAI items, respectively. With neither AI and NAI content is there a normal distribution of responses; the histogram spreads in both directions, with participants representing both ends of the spectrum. Fig. 4 with AI content reflects a positive skew, in line with our predictions about the NAI/AI distinction, while Fig. 5 reflects more of a bimodal distribution. 36.6% of participants in the AI condition and 25.6% of participants in the NAI condition chose the intra-sentential referent under 20% of the time. By contrast, only 17.1% of all participants in AI condition and 29.3% in NAI condition chose the intra-sentential referent over 80% of the time. Thus, placing the

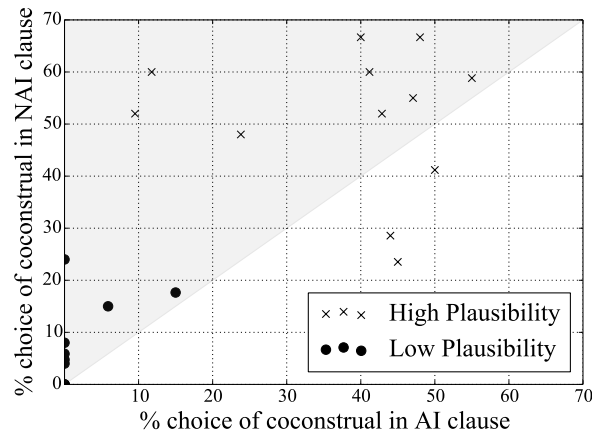


Figure 3: Percentage choice of intra-sentential referent individual test items in AI vs. NAI condition.

pronoun-name sequence in the NAI content pulled responses away from the extra-sentential referent to the intra-sentential referent, for a sizable subset of participants, raising interesting questions about the nature of the factors and the representation of the binding constraints among the population of experimental participants we tested. We return to this issue in the Discussion section.

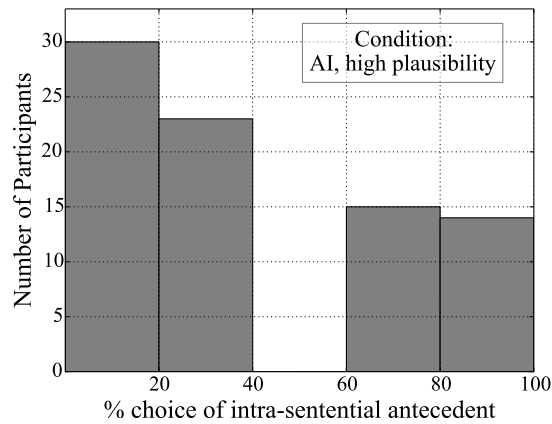


Figure 4: Distribution of participant responses reflecting choice of intra-sentential antecedent in high plausibility target items encoded in AI content.

3 Discussion

We began this paper with a question: *What are the specific pragmatic or contextual factors that influence speakers' judgments of acceptable, but structurally illicit, co-construal?* This question arose because of the striking contrast between the categorical predictions about possible coconstruals generated by Principle C and the range of acceptable coconstrual relations that encode counterexamples to Principle

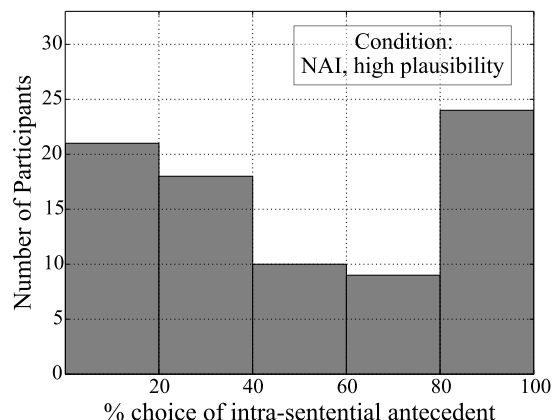


Figure 5: Distribution of participant responses reflecting choice of intra-sentential antecedent in high plausibility target items encoded in NAI content.

C, which strongly suggest that there is a trade-off between the stringent structural constraint and the influence of other factors in and outside the grammar. While earlier theoretical proposals have attempted to account for counterexamples to Principle C by appealing to context and pragmatics, these vague appeals overgenerate inferences about acceptability of coconstrual. Our goal has been to pinpoint the factors involved in licensing structurally illicit, pragmatically viable coconstrual.

In this paper, we have presented evidence for two factors that systematically affect speakers' preference: (i) the plausibility of a coconstrual relation, and (ii) the (not)-at-issue status of the content in which the pronoun-name sequence occurs. The first factor, plausibility, has independently been previously reported to exert significant influence on comprehension during sentence processing, and in previous work, we have found that it exerts an influence on coconstruals and Principle C. The current results reinforce these previous findings related to plausibility. The second factor, (N)AI status, has been shown to matter for processing and interpretation, since NAI content does not answer the QUD and projects, and while truth conditionally meaningful, does not impact processing the same way as AI content (Dillon *et al.* 2014, Dillon *et al.* 2017, Syrett & Koev 2014).

The current experiment has provided additional evidence for differential treatment of AI vs. NAI content, this time with structurally illicit coconstruals embedded in sentence-initial temporal clauses. We have demonstrated that when a pronoun-name sequence is embedded as NAI content (i.e. as part of an adjunct temporal clause), coconstrual is significantly more likely. While this finding about the differential treatment of NAI and AI content is in agreement with conclusions drawn by Dillon *et al.* (2014) and Dillon *et al.* (2017), the environments we targeted were not the same. Dillon *et al.* (2014) hypothesized that the observed distinction stems from the fact that not-at-issue content contributes a *quasi-independent* speech act from its host clause (Arnold 2007, Frazier *et al.* 2015, Potts 2005, Syrett & Koev 2014), and this quasi-independence (or *illocutionary independence*) has perceptual consequences for online sentence processing. However this claim has not been generalized to all types of NAI content, e.g., temporal clauses, which are presupposition

triggers and were the target NAI environment in the study presented here.

We wish to argue that the observed differences in processing may be due to the fact that NAI content contributes a non-negotiable update, i.e., information that is directly added to the common ground (AnderBois *et al.* 2013, Murray 2010). While main clauses introduce a proposal to update the context which can be negotiated, and then either approved or rejected, clauses that embed NAI content are *imposed* on the context in a non-negotiable way. This property is particularly strong with presupposed content (e.g., temporal clauses), which is expected to be taken for granted and not open to discussion (Fintel 2000). As a result, a structurally illicit coconstrual relation embedded in presupposed NAI content, which is not subject to negotiation, has a greater chance of being accepted *as is* by the listener than an identical one hosted in AI content, which enters the conversation as a proposal for the listener. This effect becomes particularly strong in cases where the plausibility of coconstrual is high and the relation is consistent with a pre-existing memory schema about the real world. Thus, syntax interacts with pragmatics and draws upon real world knowledge in very specific ways to give rise to coconstrual relations.

This brings us to the hypotheses we entertained: either Principle C is a hard and fast principle of the grammar, and treatment of structurally illicit coconstruals is consistent across AI and NAI content (H_0), or else Principle C is a factor that interacts with other factors within and outside of the grammar, and participants' acceptance of coconstrual is higher in NAI content, *modulo* plausibility (H_A). The evidence from our study provides support for the latter by demonstrating that both plausibility and (N)AI status of the content embedding a pronoun-name sequence contribute to participants' choice of an intrasentential referent for a pronoun.

But not all participants jump on the coconstrual bandwagon. Another important observation from our experimental results, supported by Fig. 4 and Fig. 5, is that there was significant variance in participants' answers. While some participants were primarily guided by a syntactic restriction on coconstrual and invariably chose extrasentential antecedent in the majority of cases, others were more malleable, appearing to make their selection based on a combination of factors, allowing plausibility and NAI status to override the restriction on coconstrual imposed by Principle C, thereby shifting the distribution to the right. We therefore observe that native speakers employ different strategies during sentence comprehension to resolve pronominal referential ambiguity.

This variability among our participants raises intriguing questions about these representations of Principle C in the grammars of the speakers. What are the implications of this difference in response pattern? For one group, Principle C is as Chomsky and others would have it: an inviolable principle constraining the search for an antecedent and categorically ruling out a name in the c-command domain of the pronoun. For others, Principle C is just one of several factors, and the relative weighting of these factors shifts. Does this mean that we are witnessing two different grammars, or dialects, of native speakers? It is quite possible that this is the case, and under normal circumstances, this difference would not surface. However, in these carefully controlled experiments with backwards anaphora in which we induce a Principle C effect, but manipulate the influence of other non-structural factors, we see differences in interpretive strategies we would not otherwise see. Exploring the issues that arise here is an exciting avenue for future research.

4 Conclusions

In this paper we began by sharing intuitive observations from the previous literature that Principle C, as formulated by Chomsky (1981), is insufficient in predicting when coconstrual between a pronoun and a name will be licensed. While theoretical proposals that have appealed to context and pragmatics may account for such counterexamples to the structural binding condition, they fail to allow us to make systematic predictions about speakers' judgments. This tension between under- and over-generation of coconstrual relations led us to experimentally explore the factors interacting with Principle C in resolving referential ambiguity of pronouns. We have presented experimental evidence that coconstrual between a pronoun and a name it c-commands is not solely determined by Principle C. Rather, the plausibility of coconstrual and the (N)ot-At-Issue status of the content in which the pronoun-name sequence appears also exert a robust influence on interpretation. Based on this evidence, we have argued that Principle C is not a hard and fast, inviolable principle of the grammar, but rather an indicator of which coconstrual relations are marked, and one of several factors, whose effect can be modulated. Our research therefore calls for a more thorough ongoing analysis of which other grammatical and extragrammatical factors interact with the binding constraints, and a more precise formulation of what the mechanism for establishing coconstrual is.

References

- AnderBois, S., A. Brasoveanu, & R. Henderson. 2013. At-issue proposals and appositive impositions in discourse. *Journal of Semantics* 32.93–138.
- Arnold, D. 2007. Non-restrictive relatives are not orphans. *Journal of linguistics* 43.271–309.
- Boland, J. E., M. K. Tanenhaus, S. M. Garnsey, & G. N. Carlson. 1995. Verb argument structure in parsing and interpretation: Evidence from wh-questions. *Journal of Memory and Language* 34.774–806.
- Bolinger, D. 1979. Pronouns in discourse. *Syntax and Semantics 12: Discourse and Syntax* 12.289–309.
- Chien, Y.-C., & K. Wexler. 1990. Children's knowledge of locality conditions in binding as evidence for the modularity of syntax and pragmatics. *Language acquisition* 1.225–295.
- Chomsky, N. 1981. *Lectures on government and binding*. Dordrecht: Foris.
- Clifton, C. 1993. Thematic roles in sentence parsing. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale* 47.222.
- Dillon, B., C. Clifton Jr, & L. Frazier. 2014. Pushed aside: parentheticals, memory and processing. *Language, cognition and neuroscience* 29.483–498.
- Dillon, B., C. Clifton Jr, S. Sloggett, & L. Frazier. 2017. Appositives and their aftermath: Interference depends on at-issue vs. not-at-issue status. *Journal of Memory and Language* 96.93–109.
- Evans, G. 1980. Pronouns. *Linguistic inquiry* 11.337–362.
- Fintel, K. v. 2000. What is presupposition accommodation. *Manuscript, MIT*.
- Frazier, L., B. Dillon, & C. Clifton. 2015. A note on interpreting damn expressives: Transferring the blame. *Language and Cognition* 7.291–304.
- Gobel, A. 2018. Final appositives at the right frontier: An experimental investigation of anaphoric potential. Talk presented at Sinn und Bedeutung (SuB23), Universitat Autònoma de Barcelona.
- Gor, V. 2017. Acceptable but ungrammatical in comparison: Overriding principle c effects in comparatives. In *Proceedings of 47th Annual Meeting of the North East Linguistic Society (NELS 47)*, volume 3, 45–54.

- Gor, V., & K. Syrett. 2015. Picking up after sloppy children: What pronouns reveal about children's analysis of English comparative constructions. In *Proceedings of the 39th Boston University Conference on Language Development (BUCLD 39)*, volume 1, 191–203.
- Gor, V., & K. Syrett. 2018. Principle C is just the beginning: Investigating the influence of plausibility and pronominal position on coconstrual. Ms., Rutgers University.
- Grodzinsky, Y., & T. Reinhart. 1993. The innateness of binding and coreference. *Linguistic inquiry* 24.69–101.
- Harris, J. A., & C. Potts. 2009. Perspective-shifting with appositives and expressives. *Linguistics and Philosophy* 32.523–552.
- Heim, I. 1982. The semantics of definite and indefinite NPs. *University of Massachusetts at Amherst dissertation*.
- Higginbotham, J. 1985. On semantics. *Linguistic inquiry* 16.547–593.
- Hunter, J., & N. Asher. 2016. Shapes of conversation and at-issue content. In *Semantics and Linguistic Theory*, volume 26, 1022–1042.
- Kazanina, N., E. F. Lau, M. Lieberman, M. Yoshida, & C. Phillips. 2007. The effect of syntactic constraints on the processing of backwards anaphora. *Journal of Memory and Language* 56.384–409.
- Levinson, S. C. 2000. *Presumptive meanings: The theory of generalized conversational implicature*. MIT press.
- Murray, S. E. 2010. *Evidentiality and the structure of speech acts*. Rutgers University New Brunswick, NJ dissertation.
- Pickering, M. J., & M. J. Traxler. 1998. Plausibility and recovery from garden paths: An eye-tracking study. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 24.940.
- Potts, C. 2005. *The logic of conventional implicatures*. Oxford University Press on Demand.
- Reinhart, T. 1983. Coreference and bound anaphora: A restatement of the anaphora questions. *Linguistics and Philosophy* 6.47–88.
- Roberts, C. 1996. *Information structure in discourse: Towards an integrated formal theory of pragmatics*. Ohio State University. Department of Linguistics.
- Safir, K. 2004. *The syntax of anaphora*. Oxford University Press on Demand.
- Sag, I. A. 2000. Another argument against wh-trace. *Hankamer's Webfest*. <https://linguistics.ucsc.edu/research/publications/Hankamer>
- Schlenker, P. 2010. Supplements within a unidimensional semantics I: Scope. In *Logic, language and meaning*, 74–83. Springer.
- Schwarz, F. 2007. Processing presupposed content. *Journal of Semantics* 24.373–416.
- Schwarz, F., & S. Tiemann. 2017. Presupposition projection in online processing. *Journal of Semantics* 34.61–106.
- Simons, M., J. Tonhauser, D. Beaver, & C. Roberts. 2010. What projects and why. In *Semantics and linguistic theory*, volume 20, 309–327.
- SuperLab Pro. 2015. v. 5.0 [Computer software]. Phoenix, AZ: Cedrus Corporation.
- Syrett, K., & T. Koev. 2014. Experimental evidence for the truth conditional contribution and shifting information status of appositives. *Journal of Semantics* 32.525–577.
- Tanenhaus, M. K., G. Carlson, & J. C. Trueswell. 1989. The role of thematic structures in interpretation and parsing. *Language and cognitive processes* 4.SI211–SI234.
- Tonhauser, J. 2012. Diagnosing (not-) at-issue content. *Proceedings of Semantics of Underrepresented Languages of the Americas (SULA)* 6.239–254.
- Traxler, M. J., & M. J. Pickering. 1996. Plausibility and the processing of unbounded dependencies: An eye-tracking study. *Journal of Memory and Language* 35.454–475.
- Trueswell, J. C., M. K. Tanenhaus, & C. Kello. 1993. Verb-specific constraints in sentence processing: Separating effects of lexical preference from garden-paths. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19.528–553.