Coherence-driven predictability and referential form: Evidence from English corpus data[1]

Xixian LIAO — *Universitat Pompeu Fabra*

Abstract. To refer to a discourse entity, speakers need to choose from a variety of expressions, such as a name *Donald Trump* or a pronoun *he*. While evidence has accumulated that more predictable words are more likely to be phonologically reduced, there is a long-standing debate regarding whether more reduced referring expressions (e.g., pronouns) are more frequently produced for more predictable referents. This study contributes to this debate with new evidence in two aspects. 1) Rhetorical relation-driven predictability: while the cases that have been studied in previous psycholinguistic studies have been fairly restricted to some particular verbs types, this study is, to the best of our knowledge, the first attempt of broadening the empirical base with expectation primarily driven by rhetorical relations; 2) Naturally-occurring language in corpus data: instead of using constructed language to elicit production, we make use of natural coreference chains that can be automatically retrieved from corpora developed in the field of Computational Linguistics. We found uniform pronominalization rates across rhetorical relations despite the different next-mention rates, supporting a dissociation between likelihood of next mention and pronoun production.

Keywords: predictability, rhetorical relations, corpus

1. Introduction

When referring to a previously mentioned discourse entity, speakers have a variety of expressions to use, from more explicit descriptions (*the 45th president of the United States*), or names (*Donald Trump*), to more reduced forms like pronouns (*he*). Pronouns that lack clear antecedents could lead to miscommunication. On the other hand, abusing the use of longer expressions harms efficiency. Then how do speakers choose which form to use? Intuitively, if the upcoming referent is predictable for addressees, speakers could resort to more reduced referring expressions in pursuit of efficient communication. For instance, speakers might tend to shorten the word "information" to "info" given a context like "search for more _____", where the upcoming content is predictable for comprehenders. This association between predictability and reduction has been widely suggested at the word level (e.g. Jurafsky et al., 2001).

Nonetheless, divergent claims have been made in respect of the effects of predictability at the referential level, particularly on pronoun production. Some studies found that speakers produced pronouns instead of longer expressions more frequently for more predictable referents (Arnold, 2001; Rosa and Arnold, 2017; Weatherford and Arnold, 2021), whereas other studies suggested that likelihood of a referent being the next mention is unrelated to likelihood of using a pronoun to refer to that referent (Stevenson et al., 1994; Fukumura and Van Gompel, 2010; Rohde and Kehler, 2014). This leaves the question open as to whether predictability affects the choice of referring expression.

[1]I would like to thank anonymous reviewers, audiences at SuB26 as well as my supervisors, Gemma Boleda and Laia Mayol, for their insightful comments. This study is supported by an FI-AGAUG grant (2020FI-B00575), the grant PGC2018-094029-A-I00 (AEI/FEDER, EU), and funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 715154).

 In: Daniel Gutzmann and Sophie Repp (eds.): Proceedings of Sinn und Bedeutung 26, pp. 544–556. University of Cologne.

To date, evidence on this question has largely been based on empirical experiments with carefully controlled contexts, where predictability is measured based on subjects' estimates about which referent will be mentioned next. This study, instead, measures next-mention frequency in corpus texts as an approximation for referent predictability, and contributes evidence from naturally-occurring discourse to this question.

1.1. Background

Previous studies concerning whether referent predictability influences referential choice have mostly used story continuation tasks, in which participants first need to comprehend a story context consisting of one sentence or more, and then they are asked to add a natural continuation to the story. The referent that ends up being more often re-mentioned first in continuations is considered to be more predictable than others. For instance, after viewing (1) or (2), people tend to continue with the referent which is assigned with the thematic role *goal* (e.g. *Brendan loved pie and cakes and all manner of sweet things but didn't know how to bake*).

(1) ... Lisa$_{source}$ gave the leftover pie to Brendan$_{goal}$. _____

(2) ... Marguerite$_{goal}$ caught a cold from Eduardo$_{source}$ two days before Christmas._____

To test how this expectation bias influences pronoun production while controlling for the well-known effects of grammatical roles, previous studies compare the pronominalization of the goal and that of the source when both are introduced in the same grammatical position (e.g. *Lisa* in (1) vs. *Marguerite* in (2)).

Arnold (1998) found that participants tended to continue stories with goal referents more often than source referents, and they also used more pronouns referring back to goal referents than to source. She then proposed the Expectancy Hypothesis in which referent predictability from comprehenders is closely tied to speakers' choice of referring expression, with more reduced forms used for more predictable referents (Arnold, 1998, 2001). In a similar vein, Rosa and Arnold (2017) replicated this bias of producing more pronouns for goal referents, suggesting that predictability **affects** referential choice. More recently, Medina Fetterman et al. (2022) provided cross-linguistic evidence for this view. They found that Spanish speakers tended to use more overt pronouns for goal referents than for source referents.

Intriguingly, some studies yielded conflicting results. For instance, Fukumura and Van Gompel (2010) constructed stories using another type of verb: implicit causality verbs, such as the ones in (3) and (4). In (3), *scare* is a stimulus-experiencer verb, with the subject being the stimulus and the object being the experiencer; and in (4), *fear* is an experiencer-stimulus verb with the opposite position of thematic roles. Participants in their study were found to refer to stimulus referents more often than experiencer referents for both stimulus-experiencer and experiencer-stimulus verbs, which suggested that the stimulus is referentially more predictable over the experiencer. However, no evidence was found for a stimulus bias in pronoun use. The proportion of pronouns was the same when participants referred to the stimulus in stimulus-experiencer verbs and when they referred to the experiencer in experiencer-stimulus verbs.

In addition, Fukumura and Van Gompel (2010) further examined likelihood of next-mention and likelihood of pronoun production in sentence pairs like (5) and (6), which only differ in

the connective. They found that participants tended to continue stories like (5) with the stimulus (*Gary*) when explaining the cause; whereas participants continue (6) with the experiencer (*Anna*) more frequently when talking about the result. That is, replacing the connective *because* with *so* modified the rhetorical relations between clauses and consequently, reversed the default next-mention biases elicited by implicit causality verbs. Therefore, stimulus referents are not equally predictable in (5) and (6), neither do experiencer referents. Despite this, Fukumura and Van Gompel (2010) found that participants produced pronouns referring to the stimulus (or the experiencer) in Explanation like (5) as frequently as they did in Result like (6). All this evidence led them to the conclusion that the referential choice is **unaffected** by the likelihood of next mention.

(3) Gary$_{stimulus}$ scared Anna$_{experiencer}$ after the long discussion ended in a row. This was because ...

(4) Gary$_{experiencer}$ feared Anna$_{stimulus}$ after the long discussion ended in a row. This was because ...

(5) Gary$_{stimulus}$ scared Anna$_{experiencer}$ after the long discussion ended in a row, **because** ...

(6) Gary$_{stimulus}$ scared Anna$_{experiencer}$ after the long discussion ended in a row, **so** ...

In related work, the Bayesian Probabilistic Model proposed by Kehler and Rohde (2013) also dissociates likelihood of next mention from likelihood of pronominalization. More specifically, the Bayesian model for pronouns comes in two forms. In the *weak* form, it proposes that pronoun interpretation and production are related using Bayes rules. In its *strong* form, it further posits a dissociation between pronoun production and interpretation. Crucially, the top-down expectations that comprehenders have about next mention are conditioned primarily on semantico-pragmatic factors (e.g., verb semantics, rhetorical relations), whereas the likelihood of pronominalization provides bottom-up evidence about the topichood status of referents that are specific to speakers' choice, conditioned dominantly by grammatical factors (e.g., grammatical roles) and information structure (topicality). Thus, under this view, addressees' expectations and speakers' choices do not mirror each other. Evidence supporting this view has also been found in other languages, like German (Holler and Suckow, 2016), Catalan (Mayol, 2018) and Mandarin Chinese (Zhan et al., 2016). To sum up, previous studies paint a mixed picture of the effects of predictability on pronoun production.

1.2. Current study

As shown in the previous section, scholars have long debated the relationship between referent predictability and referential choice. The current study contributes to the existing debate with new evidence in two aspects:

(1) Unlike most of the previous studies which used highly controlled contexts to elicit language production data, we examine the effects of referent predictability in naturally-occurring language from a large-scale corpus. Constructed contexts in previous psycholinguistic experiments are oftentimes confined to a "naive" world that solely involves one single event and two animate participants, such as "Gary scared Anna" or "Lisa gave the leftover pie to Brendan".

Corpus texts, by comparison, are more representative of genuine language production in more realistic scenarios. Since we do not gather estimates about the upcoming referent from comprehenders, we measure next-mention frequency in the corpus as a way to quantify referent predictability.

This operationalization is based on the assumption that comprehenders track frequency patterns in their experience and use this knowledge to predict upcoming information. Evidence for this assumption has been found for both transfer-of-possession scenarios and implicit causality ones. Specifically, Arnold (2001) examined a corpus of spoken language consisting of transcripts from the Canadian Parliament and found that speakers, in general, tended to more frequently re-mention the goal referent than the source referent in transfer-of-possession contexts. In a more recent study, Guan and Arnold (2021) examined frequency patterns in more natural implicit causality scenarios. They searched in Google contexts resembling the controlled stimuli used in previous psycholinguistic studies, namely those with two animate arguments and followed by the connective ”because”. Their results demonstrated that the implicit cause tended to be more often re-mentioned than the non-implicit cause, consistent with the next-mention biases reported by previous psycholinguistic studies. Therefore, we believe that the actual next-mention frequency in a corpus is a good approximation of referent predictability or next-mention biases measured in previous work.

(2) Both verb types and rhetorical relations (semantics of connectives) have been shown to influence referent predictability. While previous studies mostly focus on expectations induced by specific verb types and how these expectations are modulated in interaction with different rhetorical relations, the current study investigates referent predictability that is primarily driven by rhetorical relations. This is because the re-mention patterns in transfer-of-possession and implicit causality scenarios have been shown to be hard to be reproduced using corpus data (Guan and Arnold, 2021; Liao et al., 2022), They require very strict restrictions on various aspects of context, such as referent animacy, verb sense, and syntactic structure. Rhetorical relations, instead, can be expected to have a similar semantic general effect across contexts. We thus hypothesize that expectations primarily driven by rhetorical relations might be more robust and decide to look at expectations primarily driven by the discourse coherence established through rhetorical relations. We next turn to our corpus analysis with rhetorical relations.

2. Corpus analysis

2.1. Corpus

For the analysis, we assembled a dataset by extracting targeted sentences from a large-scale corpus. The corpus we used is OntoNotes (Pradhan et al., 2012), which is richly annotated with linguistic information and has been the dominant resource for anaphora research in Computational Linguistics during the last decade (see Table 1 for an annotated example).

In particular, coreference information annotated on mention spans plays an essential role in this study. As shown in (7), mentions that refer to the same entity are labeled with the same ID in each document (referents that appear just once, namely singletons, are not annotated in OntoNotes, such as *their homes* in the given example). This gold information about reference

Word	POS	Tree	Lemma	Framenet	Sense	Speaker	Named entity	predicate-argument	Coreference
A	DT	(TOP(S(NP(NP*	-	-	-	-	*	(ARG0*	(0
wildfire	NN	*)	-	-	-	-	*	*	-
in	IN	(PP*	-	-	-	-	*	*	-
California	NNP	NP*)))	-	-	-	-	(GPE)	*)	0)
forced	VBD	(VP*	force	01	1	-	*	(V*)	-
hundreds	NNS	(NP(NP*)	-	-	-	-	(CARDINAL)	(ARG1*	(1
of	IN	(PP*	-	-	-	-	*	*	-
people	NNS	(NP*)))	people	-	1	-	*	*)	1)
from	IN	(PP*	-	-	-	-	*	(ARG2*	-
their	PRP$	(NP*	-	-	-	-	*	*	(1)
homes	NNS	*)))	home	-	1	-	*	*)	-

Table 1: Multiple layers of annotation in OntoNotes.

allowed us to automatically identify which referent each mention refers to, without additional efforts on manual examination and human decision-making.

(7) A wildfire in California$_0$ forced hundreds of people$_1$ from their$_1$ homes.

As our research mainly concerns referent predictability in English, we only used the English language proportion (1.7M token), comprising data from a diverse set of genres: newswire (main source), magazine articles, broadcast news, broadcast conversations, web data, conversational speech, New Testament and Old Testament texts. The data size of each genre is shown in Table 2.

Genre of data	Size
Newswire	625k
Broadcast news	200k
Broadcast conversations	200k
Web data	300k
Telephone conversation	120k
Old Testament and New Testament	300k
Total	1745k

Table 2: Size of English data.

2.2. Methods

We made use of the morphosyntactic and semantic information annotated in OntoNotes to automatically extract rhetorical relations and label references. In the absence of discourse structure annotation, we use explicit connectives to identify relations. We left out relations that are typically signaled by ambiguous connectives. For instance, Parallel was not considered because it is typically marked by *and*, which is compatible with clauses that stand in Occasion and Contrast as well.[2] We then selected a set of connectives for the following relations: Occasion,

[2]The term "Occasion" is used in Kehler (Kehler, 2002) to refer to the relation which typically connects descriptions for temporally and spatially contiguous events. It is also called as "Narration" in Segmented Discourse Representation Theory (Asher and Lascarides, 2003) and "Sequence" in Rhetorical Structure Theory (Mann and Thompson, 1988).

Contrast, Result, Explanation, Purpose, as listed in Table 3.

	Relation	Connectives
coordinating	Occasion	(and) then, later, next, (a period of time) later/after, afterward(s), after it/that
	Result	(and) so, thus, accordingly, consequently, hence, therefore, as a result, as a consequence
	Contrast	but, on the contrary, by contrast, however
subordinating	Explanation	because
	Purpose	so that, so (annotated as the beginning of a modifier argument denoting purpose)

Table 3: Connectives used for the extraction of rhetorical relations.

We grouped the relations in accordance with the prototypical classification for coordinating and subordinating relations (Asher and Vieu, 2005). In the group of coordinating relations, there are Occasion, Contrast, and Result, which are typically realized as coordinating conjunction. We only considered inter-sentential cases in which the connective lies at the beginning of a sentence, such as (8). We left out intra-sentential cases such as "An evil spirit comes into him, **and then** he shouts". This is because sentence-internal coordinating conjunction is often accompanied by null subjects, as in "Judas went over to Jesus and then Ø kissed him". We leave the exploration of null subjects to future work (see Section 3 for discussion).

In contrast, Explanation and Purpose, relations oftentimes marked by subordinating connectives, were extracted intra-sententially like (9). These relations were extracted and analyzed separately because the distinction between coordinating and subordinating has been shown to influence mention choice, with more pronouns produced after subordinating connectives than after coordinating ones (Fukumura and Van Gompel, 2010).

(8) Occasion: <s>Judas ate the bread Jesus gave him.</*s* > <s>**Then** he immediately went out.</*s* >

(9) Explanation: <s>Some people are telling the message about Christ **because** they are jealous and bitter.</*s* >

Then for each extracted sample, the first grammatical subject after the connective was labeled as the next mention. Using the first noun phrase after the connective instead results in too much noise, such as cases in which it indicates time or location (e.g. *this week* or *school* in *at school*)

All the samples extracted for each rhetorical relation were then classified into three coreference types: the next mention is coreferential with the preceding subject, with a preceding non-subject, or with another referent that has not been mentioned in the preceding clause.

2.3. Hypotheses

For relations in the coordinating group, we put forward the following hypotheses:

(H1) Occasion will show stronger next-mention biases towards the preceding subject than Contrast and Result. The underlying rationale is the following: a) Occasion, by its definition, will display continuity in the entities which narrative sequences of events center around; and b) the canonical place to mention both the topic and the agent role is the grammatical subject. By contrast, in other relations, this subject bias could be mitigated. For instance, in a Result such as "Hurricane Maria struck Puerto Rico yesterday. As a result ...", it is quite likely that the patient role bears the consequences.

(H2) If next-mention biases exert effects on pronoun production, a comparatively higher pronominalization rate will be observed in Occasion when re-mentioning the preceding subject.

As for the relations in the subordinating group, we do not have a clear hypothesis. They were extracted for comparison so that we could examine the effects of syntactic factors on pronoun production. We expect more pronouns produced following subordinating connectives than coordinating connectives, as referents in the main clause should remain highly activated as subordinating clauses are produced (Fukumura and Van Gompel, 2010).

With respect to the effects of other grammatical/structural factors, we expect more pronouns to be used in subject coreference contexts than in non-subject coreference ones, namely, we expect a subject preference in pronoun production.

2.4. Results

With respect to H1, we looked at the distribution of different coreference types in each relation, as shown in Figure 1. Within coordinating relations, there are more contexts continuing with the subject in Occasion than in Contrast and Result ($\chi^2(2) = 20.3$, $p < .0001$; no difference between Contrast and Result $\chi^2(1) = 0.5$, $p = .5$). Within subordinating relations, there is a comparatively larger percentage of references to the subject in Explanation as opposed to Purpose ($\chi^2(1) = 4.8$, $p = .03$). In turn, the contexts of Purpose more frequently continue with non-subjects than the contexts of Explanation do ($\chi^2(1) = 21.1$, $p < .0001$), and do not show the common subject bias.

To ensure that these next-mention biases are primarily driven by rhetorical relations and not by verb semantics, we classified verbs into several semantic groups and compared the coreference distribution by verb types.[3] We zoomed in on the distribution of coreference type by verb types in each rhetorical relation. Table 4 shows that a larger proportion of subject coreference contexts in Occasion prevails in most of the verb types, except in stative predicates (e.g. mental state verbs). With eventive predicates such as transfer of possession verbs, Occasion displays prominent subject biases. Within subordinating relations, Explanation shows larger subject biases and Purpose, larger non-subject biases in almost all the verb types (see Table 5).

[3]We considered mainly the most frequent verb types that appear in the samples of rhetorical relations. The rest was collapsed into the category *others*.

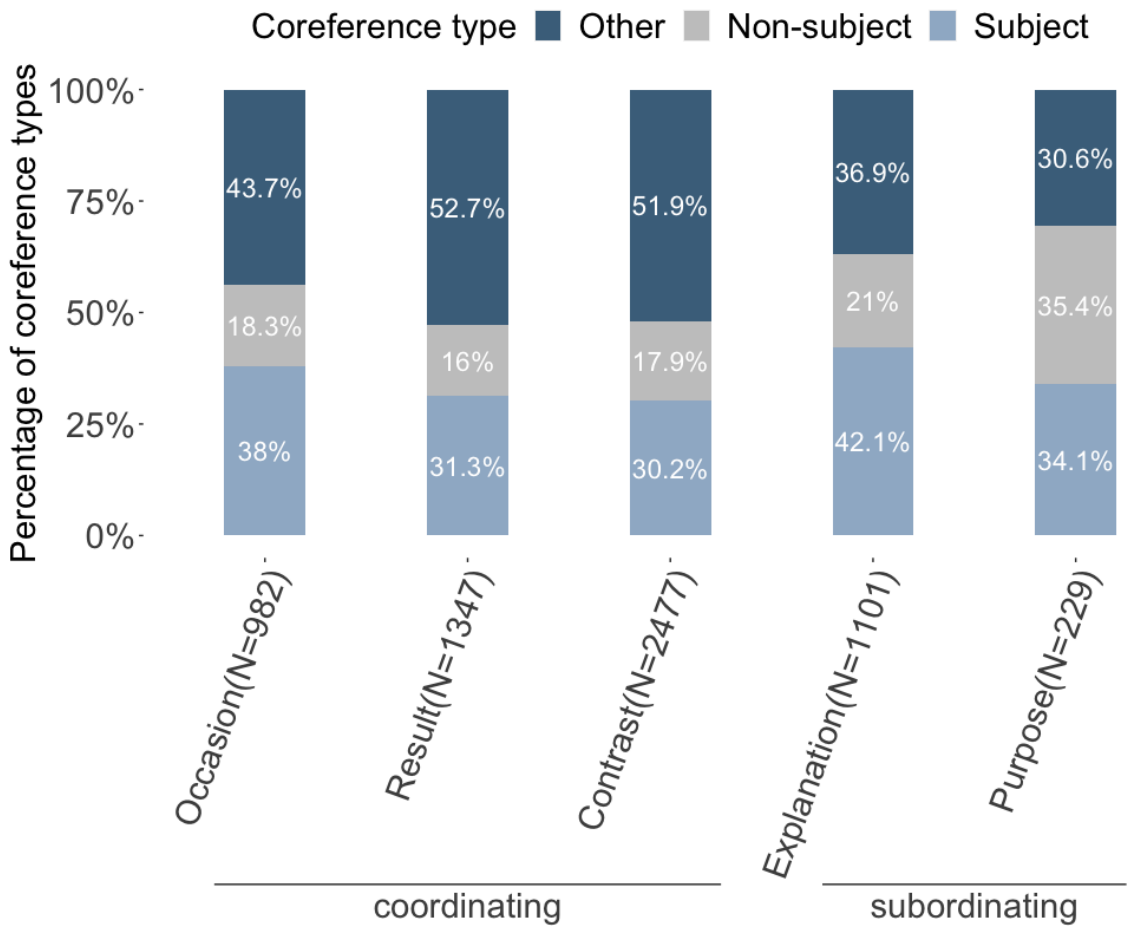

Figure 1: Coreference type by rhetorical relation in OntoNotes.

	coordinating(subject coref)		
	Occasion	Result	Contrast
mental state verbs	37.0	**52.5**	42.3
be	15.9	18.2	**21.8**
speech verbs	**34.4**	22.9	33.5
implicit causality verbs	**46.2**	29.8	29.6
transfer of possession verbs	**56.6**	21.3	27.5
do	**43.8**	28.1	36.7
others	**39.4**	38.4	31.9

Table 4: Percentage (%) of subject coreference contexts with different verb types in coordinating relations.

With respect to H2, we looked at pronominalization rates in subject coreference contexts and non-subject coreference contexts, as shown in Figure 2. Within coordinating relations, no evidence was found that more pronouns were produced in Occasion when the subject referent was re-mentioned ($\chi^2(2) >= 3.0$, $p = .2$; compare the three leftmost bars), despite the higher

	subject coref		*non-subject coref*	
	Explanation	Purpose	Explanation	Purpose
mental state	**53.4**	50.0	23.3	**25.0**
be	**32.8**	20.0	16.7	**45.0**
speech	**47.1**	20.0	13.2	**60.0**
ICV	39.0	**42.9**	30.7	**33.3**
TPV	**35.9**	20.5	43.6	**48.7**
do	**56.3**	35.0	**15.6**	10.0
others	**49.7**	39.5	16.9	**30.9**

Table 5: Percentage (%) of coreference samples with different verb types in subordinating relations.

rate of referring back to the subject in Occasion than in the other two. Within subordinating relations, pronominalization rates are similar when re-mentioning the subject referent (Fisher's exact test: $p = .23$; compare light green and yellow bars in the left); but there are more pronouns in reference to non-subject antecedents in Purpose compared to Explanation ($\chi^2(1) >= 5.9$, $p = .02$; compare the two rightmost bars).

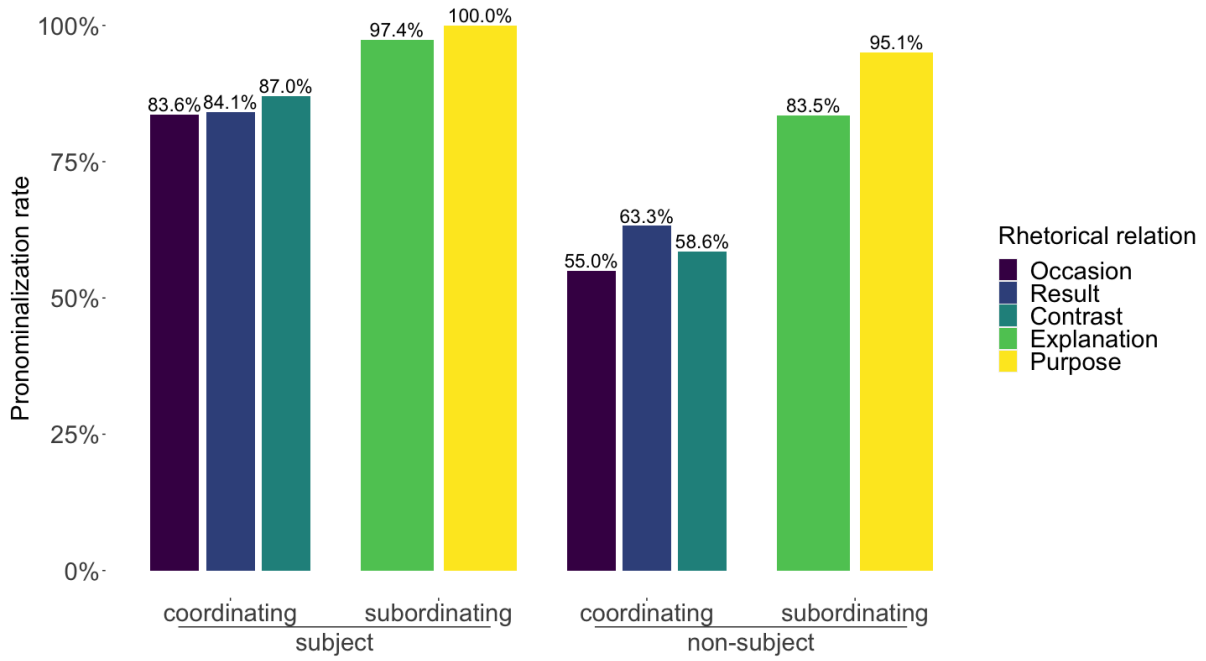

Figure 2: Pronominalization rate of next mention by rhetorical relation.

One concern is that the referents in each category may not all have the same degree of topicality. In order to check, we used prior pronominalization as a proxy for topicality. Table 6 shows the percentage of contexts where the antecedent of the next mention was already pronominal. The percentages of contexts with a pronominal subject antecedent between coordinating relations do not differ ($\chi^2(2) >= .1$, $p = .93$), neither do the percentages of contexts with a pronominal subject antecedent between subordinating relations ($\chi^2(1) >= 2.7$, $p = .10$). On the other hand,

since a larger proportion of contexts with a pronominal non-subject antecedent were found in Purpose ($\chi^2(1) >= 5.3$, $p = .02$), the non-subject biases coming from the two subordinating relations are not fully comparable.

	subject coref			*subject coref*		*non-subject coref*	
	Occasion	Result	Contrast	Explanation	Purpose	Explanation	Purpose
pro. antecedent	258	293	524	297	58	84	42
total	373	422	747	464	78	231	81
%	69.2	69.4	70.1	64.0	74.4	36.4	51.9

Table 6: Percentage (%) of samples with a pronominal antecedent.

Another concern is that we did not distinguish between different types of pronouns as antecedents, while in fact, they do not constitute a homogeneous class. First- and second-person pronouns (e.g. *I*, *you*) differ from third-person pronouns in that the former is always deictic pronouns, referring automatically to the speaker or the addressee in the utterance context, for which normally there is no other referential choice than pronouns. By contrast, for other referents, speakers/writers more plausibly have a choice between using a pronoun vs. a more explicit referring expression (e.g. name, fuller description). To ensure that the inclusion of all types of pronominal antecedents did not interfere with the validity of our previous results on pronominalization, we performed a sanity check on the subject coreference samples by excluding those with a first- or second-person pronoun subject antecedent and restricting only to contexts with a third-person pronominal or non-pronominal subject antecedent. The analysis shows that pronominalization rates in subject coreference samples remain similar across inter-sentential relations ($\chi^2(2) = 1.9$, $p = .4$) and intra-sentential relations (Fisher's exact test: $p = .4$) after excluding these potentially problematic pronouns, as presented in Table 7.

relations	subject coreference	next mention = pronoun	%
Occasion	271	223	82.2
Result	249	194	77.9
Contrast	430	350	81.4
Explanation	291	280	96.2
Purpose	42	42	100.0

Table 7: Number of subject coreference samples and pronominalization rates after excluding samples in which the subject antecedent is a first or second person pronoun.

Therefore, with respect to H2, we did not find evidence that next-mention biases exert effects on pronoun production. The uniform pronominalization pattern across coordinating relations and subordinating relations suggests, on the contrary, a dissociation between likelihood of next mention and likelihood of pronominalization.

With respect to the effects of syntactic/structural factors on pronoun production, Figure 2 shows that, as expected, there are more pronouns produced in subject coreference contexts than in non-subject conference ones (coordinating: $\chi^2(1) >= 206.9$, $p < .0001$; subordinating: $\chi^2(1) >= 40.4$, $p < .0001$). In addition, subordinating relations in general obtain higher

pronominalization rates than coordinating ones, especially when referring to non-subject referents (subject: $\chi^2(1) >= 59.7$, $p < .0001$; non-subject: $\chi^2(1) >= 76.2$, $p < .0001$). These patterns are aligned with the findings from Fukumura and Van Gompel (2010).

3. Conclusion and discussion

This study sets a precedent of using large-scale corpora developed in Computational Linguistics to investigate the relationship between referent predictability and pronoun production. Previous evidence in this field has largely been based on empirical experiments with carefully controlled contexts. Our study manifests the feasibility of implementing natural language processing techniques to extract restrictive contexts from corpora that are richly annotated with linguistic information.

We provide complementary evidence from corpus data supporting the dissociation between likelihood of next-mention and likelihood of pronoun production posited by the *strong* Bayesian Model (Kehler and Rohde, 2013), particularly in that likelihood of next-mention varies depending on rhetorical relations, whereas likelihood of pronoun production does not. This is argued for in a two-step fashion:

a) We firstly provided evidence for the hypothesized next-mention frequency pattern posited in the section 2.3. Specifically, we found a higher frequency of re-mentioning the preceding subject in Occasion over other coordinating relations (Result and Contrast) while controlling for the effects from verb semantics. It is noteworthy that while Occasion exhibits a stronger subject bias with most of the verb types, it appears to be allergic to stative predicates (mental state verbs and *be* verb in Table 4). This asymmetry is in line with Altshuler (2016)'s prediction that the continuity displayed by Occasion is sensitive to the aspectual distinction between eventive and stative in virtue of the fact that the stative descriptions, unlike the eventive ones, do not move the narrative time forward.

b) After establishing that Occasion is more likely to continue with the grammatical subject (i.e. subject referents are more predictable in Occasion), we proceeded to examine whether these contrasting likelihoods influence pronoun production. The results show that pronoun production is subject to grammatical/structural factors (referring to the preceding subject vs. referring to a preceding non-subject; relations presented in coordinating structures vs. relations in subordinating structures). On the other hand, we observed uniform pronominalization rates across rhetorical relations despite the different likelihoods of next-mention. This leads us to the conclusion that predictability, disentangled from structural factors and primarily driven by rhetorical relations, does not affect pronoun production.

Yet our conclusion has been drawn based on next-mention rates and pronominalization rates of referents in the subject position, for which the effects of predictability on pronoun use might be much weaker than that of subjecthood. Existing evidence for the claim that predictability affects pronoun production mostly comes from the object position (Arnold, 2001; Weatherford and Arnold, 2021). This opens the question of the scope of the effects of predictability. It could be possible that predictability typically affects the use of pronouns for referents previously mentioned in non-subject positions. The evidence from the present study is insufficient for us

to conclude on this point. Although we did observe a higher rate of re-mentioning non-subjects and a higher pronominalization rate for non-subjects in Purpose as opposed to Explanation, numbers in these two relations are not fully comparable given the higher prior pronominalization rate in Purpose. Additional research is required to understand better when predictability has effects on pronoun production and when it does not.

In future work, we would also like to see how our results extend to the case of zero pronouns, such as elliptical zeros in coordinate conjunctions (*John panicked and then Ø ran*) – arguably, an even more extreme case of reduction than pronouns. Previous studies have integrated zeros in different ways. In some studies (Rosa, 2015; Arnold and Nozari, 2017), zeros were considered on a par with pronouns, since they were considered to be used in similar discourse contexts. Instead, Kibrik et al. (2016) argue that syntactically induced zeros should not be treated as a discourse-based referential choice in the same way that third-person pronouns or noun phrases are. This question is critical because we would have a different picture if instances containing zeros were also considered in our analyses. Our study is restricted to inter-sentential instances of rhetorical relations where an explicit subject is usually required to be present. In *intra*-sentential contexts, we would expect zeros to be frequently used in relations like Narration. For instance, Arnold and Nozari (2017) found that speakers were more likely to use pronouns or zeros when they marked discourse connectivity with words like *and* or *then*. These zeros in coordinate constructions would increase the rate of subject re-mentioning and the rate of using a more reduced expression at the same time, and presumably result in a correlation between the two. Future research will need to take a closer look at the referential role of zeros, and obtain more insights by analyzing the effects of referent predictability on a broader range of reduction phenomena in reference production.

To conclude, with the current work, we have contributed complementary evidence from corpus data to the ongoing debate about the relationship between referent predictability and pronoun production. Our results suggest that predictability does not affect the use of pronouns. We hope that our corpus study will spark interest in the use of resources developed in the field of Computational Linguistics to investigate theoretical questions.

References

Altshuler, D. (2016). *Events, states and times: An essay on narrative discourse in English*. Berlin: de Gruyter.

Arnold, J. E. (1998). *Reference form and discourse patterns*. Ph. D. thesis, Stanford University.

Arnold, J. E. (2001). The effect of thematic roles on pronoun use and frequency of reference continuation. *Discourse processes 31*(2), 137–162.

Arnold, J. E. and N. Nozari (2017). The effects of utterance timing and stimulation of left prefrontal cortex on the production of referential expressions. *Cognition 160*, 127–144.

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

Asher, N. and L. Vieu (2005). Subordinating and coordinating discourse relations. *Lingua 115*(4), 591–610.

Fukumura, K. and R. P. Van Gompel (2010). Choosing anaphoric expressions: Do people take

into account likelihood of reference? *Journal of Memory and Language 62*(1), 52–66.

Guan, S. and J. E. Arnold (2021). The predictability of implicit causes: Testing frequency and topicality explanations. *Discourse Processes*, 1–27.

Holler, A. and K. Suckow (2016). How clausal linking affects noun phrase salience in pronoun resolution. In A. Holler and K. Suckow (Eds.), *Empirical perspectives on anaphora Resolution*, pp. 61–85. Berlin: de Gruyter.

Jurafsky, D., A. Bell, M. Gregory, and W. D. Raymond (2001). Probabilistic relations between words: Evidence from reduction in lexical production. *Typological studies in language 45*, 229–254.

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

Kehler, A. and H. Rohde (2013). A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation. *Theoretical Linguistics 39*(1-2), 1–37.

Kibrik, A. A., M. V. Khudyakova, G. B. Dobrov, A. Linnik, and D. A. Zalmanov (2016). Referential choice: Predictability and its limits. *Frontiers in psychology 7*, 1429.

Liao, X., G. Boleda, and L. Mayol (2022). Pronominalization is unrelated to referent predictability: A study on English corpus data [manuscript submitted for publication]. *Universitat Pompeu Fabra*.

Mann, W. C. and S. A. Thompson (1988). Rhetorical structure theory: Toward a functional theory of text organization. *Text 8*(3), 243–281.

Mayol, L. (2018). Asymmetries between interpretation and production in catalan pronouns. *Dialogue & Discourse 9*(2), 1–34.

Medina Fetterman, A. M., N. N. Vazquez, and J. E. Arnold (2022). The effects of semantic role predictability on the production of overt pronouns in Spanish. *Journal of Psycholinguistic Research*, 1–26.

Pradhan, S., A. Moschitti, N. Xue, O. Uryupina, and Y. Zhang (2012). CoNLL-2012 shared task: Modeling multilingual unrestricted coreference in OntoNotes. In *Joint Conference on EMNLP and CoNLL - Shared Task*, Jeju Island, Korea, pp. 1–40. Association for Computational Linguistics.

Rohde, H. and A. Kehler (2014). Grammatical and information-structural influences on pronoun production. *Language, Cognition and Neuroscience 29*(8), 912–927.

Rosa, E. C. (2015). *Semantic role predictability affects referential form*. Ph. D. thesis, The University of North Carolina at Chapel Hill.

Rosa, E. C. and J. E. Arnold (2017). Predictability affects production: Thematic roles can affect reference form selection. *Journal of Memory and Language 94*, 43–60.

Stevenson, R. J., R. A. Crawley, and D. Kleinman (1994). Thematic roles, focus and the representation of events. *Language and Cognitive Processes 9*(4), 519–548.

Weatherford, K. C. and J. E. Arnold (2021). Semantic predictability of implicit causality can affect referential form choice. *Cognition 214*, 104759.

Zhan, M., R. P. Levy, and A. Kehler (2016). Bayesian pronoun interpretation in Mandarin Chinese. In A. Papafragou, D. Grodner, D. Mirmar, and J. C. Trueswell (Eds.), *Proceedings of the 38th annual conference of the cognitive science society*, pp. 2393–2398.