

Asymmetries between interpretation and production in Catalan pronouns

Laia Mayol

Universitat Pompeu Fabra

LAIA.MAYOL@UPF.EDU

Editor: Vera Demberg

Submitted 03/18; Accepted 11/18; Published online 12/18

Abstract

The literature on Romance null-subject languages has often postulated a division of labor between Null and Overt pronouns: Nulls prefer to retrieve an antecedent in subject position, whereas Overts prefer an antecedent in a lower syntactic position (Carminati, 2002). However, recent research on English pronouns (Rohde and Kehler, 2014) has shown grammatical function alone cannot explain pronoun interpretation. According to this model, pronoun interpretation and production are sensitive to different sets of factors and, instead of being mirror images of each other, are related probabilistically in a Bayesian fashion. This paper tests this model with Catalan data from two discourse-completion experiments to study the structural, grammatical factors and the semantico-pragmatic factors that affect the interpretation and production of Null and Overt pronouns. Our main result is that both Null and Overt pronouns present asymmetries regarding their interpretation and production: (1) the production of Null pronouns is affected mainly by grammatical factors (they are more likely to be used to refer to subjects), but their interpretation is also influenced by semantico-pragmatic factors (rhetorical relations and the verb's lexical semantics), and (2) while Overt pronouns have a strong interpretation bias towards the object, they are not the preferred form to refer to the object.

Keywords: pronouns, anaphora, null pronouns, overt pronouns, reference, rhetorical relations, Catalan

1. Introduction

The production and interpretation of referring expressions is essential for successful communication: speakers need to choose one particular way to refer to the entities they want to talk about; hearers need to assign a discourse referent to the referring expressions in a discourse. A long-standing idea is that more reduced anaphoric expressions (such as pronouns) tend to be used for more accessible referents, while more complex expressions will be used to refer to less accessible referents (Ariel, 1990; Givón, 1983; Gundel et al., 1993). This raises the question of which specific factors contribute to accessibility, and whether the same factors drive both production and interpretation.

A set of factors that are well-known to influence pronoun interpretations are structural, grammatical factors. In particular, subjects make particularly good antecedents (Crawley et al., 1990; Arnold, 1999), as shown by (1), from Kehler et al. (2008). Although the semantic content of both sentences is identical, the preferred antecedent of the pronoun changes: the preferred interpretation of the pronoun is that it refers to the antecedent in subject position of the previous clause.

- (1) a. Bush narrowly defeated Kerry, and special interests promptly began lobbying **him**.
[him=Bush]
b. Kerry was narrowly defeated by Bush, and special interests promptly began lobbying **him**. [him=Kerry]

However, in other occasions, structural factors are overridden by semantico-pragmatic factors and the subject preference is not observed, as illustrated by the minimal pairs in (2). While we still observe a subject preference for (2-a), it disappears in (2-b). This difference is partly due to the verb type: both *apologize* and *scold* are so-called implicit causality verbs (Caramazza et al., 1977; Brown and Fish, 1983; Stevenson et al., 1994) which attribute the cause of the event to one of the two referents (the subject in the case of *apologize*, the object in the case of *scold*). This is precisely the referent to which the pronoun is more likely to refer. These semantic biases are in force in restricted discourse conditions, in particular when the second sentence gives a cause for the event in the first sentence (i.e. when there is a rhetorical relation of Explanation between the two sentences). A different rhetorical relation may create a different bias, as shown in (2-c), in which the event of the second sentence follows the event of the first sentence and a narration is established.

- (2) a. Mike apologized to Joe because he was late.[he=Mike]
b. Mike scolded Joe because he was late. [he=Joe]
c. Mike scolded Joe and then he left. [he=Mike]

Rohde and Kehler (2014) present a model which aims to capture both structural and (non-structural) semantico-pragmatic factors. In their model, semantico-pragmatic factors (including verb type or rhetorical relations) affect only pronoun interpretation, while structural factors affect both pronoun interpretation and production. Thus interpretation and production are fundamentally asymmetrical and this asymmetry can be modeled using Bayes' Rules (see Section 2.3 for a more full-fledged explanation).

While some research has been devoted to study how structural and semantico-pragmatic factors interact in languages such as English, Romance languages such as Catalan, the language under study in this paper, have an additional aspect to consider given their dual system of pronominal forms: these languages have both overt and null pronouns (Rigau, 1986). Moreover, the interaction of structural and semantico-pragmatic factors has not been thoroughly studied for either of the two types of pronouns. This paper aims to fill this gap and answer the following questions: 1) How do the structural and semantico-pragmatic factors affect both null and overt pronouns (Nulls and Overts, henceforth) in Catalan? 2) Can the behavior of Nulls and Overts be modeled by a probabilistic model such as the one proposed by Rohde and Kehler (2014)?

This paper is structured as follows. Section 2 presents the relevant background on pronoun interpretation and production. Sections 3 and 4 present two discourse-completion studies whose aim is to study the production and interpretation of pronouns in Catalan. The quantitative data that emerges from these experiments are useful to describe how the two pronouns are used and understood and, moreover, gives us the necessary information to put the probabilistic model to a test. Finally, section 5 ends the article with some conclusions.

2. Background

2.1 Effect of grammatical factors on pronoun interpretation

As mentioned in the introduction, it has recurrently been proposed in the literature that grammatical factors affect the interpretation of pronouns: that is, the grammatical function of a referent can make it a more or less likely antecedent for a pronoun. Particularly well-studied is the Subject-Assignment Strategy (Crawley et al., 1990), which postulates that a referent mentioned in subject position becomes a more likely antecedent for a pronoun than referents in other syntactic positions.

The literature on Romance language has used this idea to explain the interpretation of Nulls and Overts, and many studies have postulated a division of labor between both types of forms. For instance, in several corpus studies, it has been found that Nulls and Overts carry different biases: a null pronoun prefers a subject antecedent, and an overt pronoun a non-subject antecedent (Cameron, 1992; Silva-Corvalán, 1977). Carminati (2002) called this asymmetry the Position of Antecedent Hypothesis (PAH), as defined in (3).

- (3) Position of Antecedent Hypothesis: Nulls prefer to retrieve an antecedent in the highest subject position, whereas Overts prefer an antecedent in a lower syntactic position.

Carminati found support for the PAH in questionnaire and reading-time experiments for Italian. In her Experiment 1, participants were presented with two-sentence discourses, in which the second sentence contained a potentially ambiguous Null (see (4-a)) or Overt (see (4-b)). Then they had to answer a question about the referent of such pronoun (4-c). She found clear opposed biases between Null and Overts: Nulls pronouns preferred to retrieve the antecedent in subject position ('Marta' in the examples), and Overts the antecedents in object position ('Piera' in the examples).

- (4) a. Marta scriveva frequentemente a Piera quando \emptyset era negli Stati Uniti.
 "Marta wrote frequently to Piera when \emptyset was in the United States."
 b. Marta scriveva frequentemente a Piera quando lei era negli Stati Uniti.
 "Marta wrote frequently to Piera when she was in the United States."
 c. Who was in the States?

Support for the PAH has also been found for other Romance languages such as Spanish (Alonso-Ovalle et al., 2002; Keating et al., 2016) and Catalan (Mayol and Clark, 2010).¹ Other researchers have proposed that the asymmetry should not be cast in syntactic terms, but in information-structural terms, such that Nulls tend to indicate topic continuity, and Overts indicate topic change (Vallduví (1992) for Catalan, Samek-Lodovici (1996) and DiEugenio (1998) for Italian). Both theories make the same predictions in most cases, given that the referent in the subject position is usually also the topic, but would differ in the cases in which the topic is realized by another grammatical function. In some previous work (Mayol, 2010*b*), I showed that the syntactic preferences of Nulls are not a byproduct of their informational structure preferences, but that the two levels interact and the PAH needs to be refined in order to capture the pronouns' preferences. Nulls have a simple preference

1. As an anonymous reviewer points out, the support for the PAH in Alonso-Ovalle et al. (2002) is only partial: while Nulls displayed a strong subject preference, Overts did not have a strong bias to either antecedent.

for subject antecedents, regardless of whether they are topics or not. Overts have a more complex preference for low-salience (non-subject, non-topics) antecedents.²

Another important difference that has been pointed out in the literature between the two types of pronouns in Romance languages is that Overts often carry a contrastive flavor, absent in Nulls. This has led some authors to treat Overts as a counterpart to stressed pronouns in languages like English, and to account for their use as ways to encode Focus (Luján, 1985, 1999) or Contrastive Topics (Mayol, 2010a).

Having briefly examined the role of grammatical factors in pronominal references, let us now turn to the less-studied effect of semantico-pragmatic factors.

2.2 Effect of semantico-pragmatic on the interpretation of pronouns

In the last decade it has become clear that grammatical factors alone cannot explain pronoun interpretation. Take for instance, example (5) (which extends example (1) discussed above). (5-a) and (5-b), as mentioned, have been used to argue for the Subject-Assignment Strategy, since the use of the active or passive voice alters the referent assignment. However, why is this strategy not active in (5-c)? In order to account for cases like this, it has been proposed that pronouns refer to an antecedent with the same grammatical function (Smyth, 1994): a subject pronoun would be biased to a subject antecedent, and an object pronoun to an object antecedent. Unfortunately, this hypothesis cannot account for (5-d), in which a subject pronoun is interpreted as referring to the previous object.

- (5)
- a. Bush narrowly defeated Kerry, and special interests promptly began lobbying him. [him=Bush]
 - b. Kerry was narrowly defeated by Bush, and special interests promptly began lobbying him. [him=Kerry]
 - c. Bush narrowly defeated Kerry, and Romney absolutely trounced him. [him=Kerry]
 - d. Bush narrowly defeated Kerry and he quickly demanded a recount. [him=Kerry]

Kehler (2002) notices that the rhetorical relations involved in (5) are not the same and argues that rhetorical relations affect pronoun interpretation. Examples (5-a) and (5-b) are examples of Occasions. The rhetorical relation of Occasion allows a speaker to signal a narration event, in which a set of events are temporally ordered: the initial state of the second sentence is equated with the final state of the first one. Thus, in an Occasion relation, the most salient referents at the end state of the first utterance become likely antecedents for upcoming pronouns. The most salient referent will be the topic of the sentence (the referent about which the sentence is about), which typically corresponds to the subject. Hence, many pronouns appearing in Occasion relations are subject-biased. That is, the discourse in (5-a) is about what Bush did and so the pronoun refers to him. Similarly (5-b) is about what Kerry did, so the pronoun refers to him. In contrast, the discourses in (5-c) and (5-d) are not Occasions. (5-c) is a case of Parallel, in which commonalities and differences between entities are highlighted: both sentences express a similar eventuality in which either Bush or Romney defeat Kerry and, thus, the pronoun is naturally assigned to the previous object. Finally, (5-d) is a Result and, therefore, a causal link is established between the two sentences, so that the second sentence expresses a consequence of the first. Our world-knowledge pushes the antecedent

2. Since the two types of categories (syntactic and informational) largely overlap, this work will focus on the study of syntactic functions.

assignment to the object, given that it is the one who loses an election the one who typically demands a recount.³

The conclusion from the previous discussion is that rhetorical relations greatly affect pronoun interpretation and come with their own biases. Moreover, rhetorical relations often interact with other semantic factors, such as verb type. This was partially illustrated in example (2) with implicit causality verbs: some verbs carry their own biases in cases of Explanation (when the event in the second utterance explains why the event of the first utterance occurred). Another well-studied semantic class of verbs is the transfer of possession verbs (TPV, henceforth), such as ‘give’, which express the event of an object being transferred from a Source argument to a Goal argument (Stevenson et al., 1994; Arnold, 2001; Kehler et al., 2008; Rosa and Arnold, 2017). Both the Source and the Goal can function as subjects or objects depending on the verb, as illustrated in (6). Previous research (Stevenson et al., 1994; Arnold, 2001) used the discourse-completion paradigm and showed that Goal referents are more accessible, particularly when they are in subject position (6-a), while the bias is milder when the Goal is not the subject (6-b).

- (6) a. Bob received a book from Peter. He ...
 b. Peter handed a book to Bob. He ...

Kehler et al. (2008) replicated this finding: in items in which the Goal was not the subject (as in (7-b)), the pronoun did not show a clear interpretation bias (half of the time participants used the pronoun to refer to Peter; the other half to refer to Bob). However, once the data was grouped by rhetorical relations very clear patterns emerged: Occasions and Results were clearly biased to the non-subject, while Elaborations (a relation in which both sentences describe the same event) and Explanations were clearly biased towards the subject. Why should this be so? As mentioned, pronouns occurring in an Occasion relation are biased towards the most salient referent at the end state of the event in the first utterance, because in an Occasion the events are temporally ordered: the end state of the first event is followed by another event, so whatever referent is salient at the end of the first event is likely to be picked up again in subsequent events. While this referent is usually the subject, Kehler et al. (2008), propose that in a context with a TPV, the Goal argument is the most salient one, since it is the one in focus at the end of the event. That is, if the speaker is narrating a series of events and she mentions that Peter handed a book to Bob, a very likely way to continue the discourse would be to explain what Bob did with the book: by the end of the first sentence, the most salient referent is Bob. The strong non-subject bias of the Result relation can be explained in similar terms: if the speaker wants to talk about the consequences of Peter giving a book to Bob, she will probably continue talking about Bob. In contrast, other relations do not focus on the end state of the event in the first utterance: for instance, Elaboration and Explanations. These relations will be biased towards the subject. For instance, if the speaker wants to elaborate on how or why the handing took place, she will probably continue talking about Peter and not about Bob. Thus, we can conclude that rhetorical relations are biased towards a particular discourse referent: in the case of a TPV context, Elaborations and Explanations are biased towards the subject, and Occasions and Results are biased towards the non-subject.

Another semantico-pragmatic factor that has been shown to influence pronoun interpretation is the aforementioned verbs’ implicit causality. Implicit causality verbs (ICVs) have been used

3. An anonymous reviewer points out that, although the pronoun can be coerced to refer to the object, s/he finds (5-d) fairly unnatural. See section 2.3 for more discussion on the interaction between structural and pragmatic factors.

recurrently to study the role of semantic factors in the interpretation of pronouns (Caramazza et al., 1977; Brown and Fish, 1983; Fukumura and Van Gompel, 2010; Rohde and Kehler, 2014). ICVs attribute the cause of the event they denote either to the subject (ICV1s, henceforth) or to the object (ICV2s, henceforth)⁴, and, therefore, impose different biases on the pronouns that follow them. For instance, the verb ‘surprise’ is an ICV1 and, therefore, (7-a) will more likely be interpreted as conveying that it was Anne, the subject, who aced the exam and that this surprised Mary. In contrast, ‘congratulate’ is an ICV2, and (7-b) will more likely be interpreted as conveying that it was Mary, the object, who aced the exam and that is why Anne congratulated her.

- (7) a. Anne surprised Mary because she aced the math exam.
b. Anne congratulated Mary because she aced the math exam.

2.3 Relationship between interpretation and production

There is a current debate in the literature about the relationship between interpretation and production, and particularly about whether interpretation and production are affected by the same set of factors. In particular, the discussion revolves around whether the semantico-pragmatic factors that affect pronoun interpretation (see previous section) also shape production. On the one hand, Arnold (1998, 2001) has argued that both interpretation and production are affected by semantico-pragmatic factors. This idea is captured by her Expectancy Hypothesis, according to which more accessible entities are more likely to be mentioned again in discourse, which in turn increases their probability of being pronominalized. This hypothesis has been tested with transfer of possession verbs (Arnold, 2001; Rosa and Arnold, 2017): in these studies, participants used pronouns more often to refer to the Goal than to refer to the Source. The effect was particularly strong for non-subjects; in contrast, it was either weak or non-existent for subjects (in Arnold (2001) and in experiments 2 and 3 in Rosa and Arnold (2017)). On the other hand, other studies have found that semantic-pragmatic factors did not affect the choice of anaphoric form. Fukumura and van Gompel (2010) used implicit causality verbs and found that the rate of pronominalization towards either of the arguments was constant regardless of whether the verb was biased to the subject or the object. A similar finding is reported in Rohde and Kehler (2014).

Rohde and Kehler (Rohde and Kehler, 2014; Kehler and Rohde, 2013) have argued for a model in which pronoun production and interpretation are not mirror images of each other, but are sensitive to different factors. This could explain an interesting asymmetry found in the discourse-completion study by Stevenson et al. (1994). As mentioned, when participants were forced to continue with a pronoun, as in (8-a), it was not clearly biased towards either of the antecedents. In contrast, when they could choose which referring expression to use to continue the discourse, as in (8-b), participants mostly used a pronoun to refer to the subject and a name to refer to the object.

- (8) a. Peter handed a book to Bob. He ...
b. Peter handed a book to Bob.

If speakers mostly use a pronoun to refer to the subject, why doesn’t the pronoun display a stronger subject preference? The answer in Rohde and Kehler (2014) is that pronoun production and inter-

4. Fukumura and Van Gompel (2010) use the terminology ‘stimulus-experiencer (SE)’ and ‘experiencer-stimulus (ES)’. I will use the more neutral ICV1 and ICV2, given that some of the arguments of the verbs I will use are not either stimuli or experiencers, but rather agents and themes, like for instance the arguments of the verb *congratulate* in (7b).

pretation are not mirror images of each other, but are sensitive to different factors. In particular, while pronoun interpretation is affected by both structural and semantico-pragmatic factors, pronoun production is not affected by the latter. This explains the asymmetry in the two conditions of Stevenson et al. (1994): although speakers mostly use pronouns to refer to the previous subject, hearers can easily assign a pronoun to the non-subject if the semantico-pragmatic factors (i.e. the context, rhetorical relation or the verb’s implicit causality, for instance) push in that direction.

Rohde and Kehler (2014) show that, although this asymmetry between production and interpretation may not be intuitive, it is, in fact, expected if production and interpretation are related probabilistically, through Bayes’ Rule, illustrated in (9).

(9)

$$P(\textit{referent} \mid \textit{pronoun}) = \frac{P(\textit{pronoun} \mid \textit{referent}) P(\textit{referent})}{P(\textit{pronoun})}$$

The term in the left-hand side, $P(\textit{referent} \mid \textit{pronoun})$, is the probability to refer to a particular referent given that a pronoun has been used. It, thus, represents the point of view of the interpreter: he has heard a pronoun and needs to assign it a referent. The term $P(\textit{pronoun} \mid \textit{referent})$, in contrast, represents the point of view of the speaker: given that she wants to refer to a particular referent, should she use a pronoun? The crucial point is that these two probabilities are not mirror images of each other, given that a third probability plays a role: $P(\textit{referent})$. This is the probability that a particular referent will be mentioned regardless of the form.⁵

The production bias, $P(\textit{pronoun} \mid \textit{referent})$, is basically affected only by grammatical factors: i.e. there is a production bias towards pronominalizing the previous subject referent. In contrast, the interpretation bias is affected both by the grammatical factors that affect the production bias and the factors that affect whether a particular referent will be mentioned next (i.e. $P(\textit{referent})$). The latter factors are predicted not to affect production.

Rohde and Kehler (2014) showed how this model can explain the behavior of pronouns in English in a variety of contexts and that it is superior to other models: (i) the one they call Expectancy model (after the Expectancy Hypothesis in Arnold (1999)), which assumes that the interpretation bias of a pronoun towards a referent (that is, $P(\textit{referent} \mid \textit{pronoun})$) is the probability that the referent is mentioned again ($P(\textit{referent})$) and (ii) the one they call Mirror Model, which assumes that the bias of a pronoun towards a referent is the probability that such referent will be pronominalized ($P(\textit{pronoun} \mid \textit{referent})$).

One of the goals of this paper is to examine whether this model can accurately account for the behavior of Nulls and Overts in Catalan. To this end, two discourse-completion studies were carried out. In both experiments, the context sentence contains a verb which manipulates the context in a particular way: Experiment 1 uses transfer of possession verbs and Experiment 2 implicit causality verbs. The reasons for this choice were two-fold: first, each verb type is expected to interact with rhetorical relations in a particular way and, second, since they have been studied extensively, we will be able to compare our results with previous research.

5. The denominator, $P(\textit{pronoun})$, is the probability that a pronoun is used and can be computed by summing the terms in the numerator by all possible referents (subject and object). This term contributes a constant factor which normalizes the values so that they are probabilities and, therefore, sum to 1. It will not be discussed further in the rest of the paper.

3. Experiment 1: transfer of possession verbs

Experiment 1 is a discourse-completion study in which the context sentence used TPVs, which, as mentioned, usually trigger many continuations about the Goal argument. In particular, only verbs that locate the Goal in indirect object⁶ position are used so that the typical biases are reversed and we obtain a context biased towards the object. This atypical context is useful to study the interaction between grammatical and semantico-pragmatic factors.

3.1 Methods

Materials

This experiment replicates in Catalan the aforementioned experiments in Stevenson et al. (1994) and Kehler et al. (2008), including a condition for the overt pronoun. Thus, the experiment has three conditions, as can be seen in (10):

- (10) a. Condition 1: Null prompt.
El Pere li va passar un llibre al Robert. \emptyset ...
Peter passed a book to Robert. \emptyset ...
- b. Condition 2: Overt prompt.
El Pere li va passar un llibre al Robert. Ell ...
Peter passed a book to Robert. He ...
- c. Condition 3: Free prompt.
El Pere li va passar un llibre al Robert. ...
Peter passed a book to Robert. ...

The first two conditions will provide interpretation data for Nulls and Overts respectively: participants will need to interpret the pronouns and provide a completion coherent with their interpretation. The last condition ('Free condition', henceforth) will provide production data, since participants are free to choose whatever form they deem appropriate for the subject position of the completion.

The experiment included 18 critical items, all containing a TPV. In each sentence, two referents of the same gender were mentioned; half of the items contained two male proper names and the other half two female proper names. The source of the transfer event always appeared in subject position, the goal appeared as the indirect object, and the transferred object as the direct object. Three lists were built, so that each participant only saw each item in one of the conditions. Each list also contained 18 fillers, with non-TPVs and with prompts containing a connector or a temporal expression. The full list of critical items with their English translations can be seen in Appendix 1.

Based on the previous literature, we formulate the following hypotheses:

1. Nulls are expected to receive more subject interpretations than Overts.
2. Type of prompt is expected to affect rhetorical relations. If indeed Nulls trigger more subject interpretations, we expect more subject-biased rhetorical relations (Explanation and Elaboration) in the Null condition than in the Overt condition.
3. Pronoun interpretation is expected to be affected by rhetorical relations. We expect Occasions and Results to be highly object-biased and Elaboration and Explanation to be subject-biased.

6. In this section, whenever I use the label 'object' I am referring to the indirect object.

4. Pronoun production is not expected to be affected by rhetorical relations. In the Free condition, the rate of Nulls for a particular referent is expected to be consistent across different rhetorical relations.

Procedure

The data was collected in an online survey site (<https://app.surveymzmo.eu>). First, the participants read the instructions in which the procedure was explained. Then, each item was presented once at a time and the participants were asked to write a single, complete sentence as a completion. They were told to write the first completion that came to mind, avoiding humor.

One of the difficulties of working with a Null prompt is that participants needed to understand what a Null is. In order to achieve this, the instructions contained a brief informal explanation. The instructions explained the difference between a null subject and an overt subject (expressed by means of a pronoun or another phrase) and contained examples of each type of subject. There were several examples in which the Null referred both to the previous subject or to the previous object, in order not to bias the participants.

Participants

Ninety participants took part in the experiment. They were all native speakers of Catalan and students at the Universitat Pompeu Fabra. They were entered in a raffle to win a gift certificate.

3.2 Results

A total of 1620 completions (18 items * 90 participants) were collected. Two judges, the author of the paper and a linguistics graduate student at UPF, coded the antecedent of the subject of the completion into one of the following categories: Subject (if it referred to the Source), Object (if it referred to the Goal), Joint (plural reference to both the Subject and the Object), Other (if it referred to the transferred object or to something else) or Unsure (if the pronoun could be understood as referring to more than one referent).

The results reported in this paper concern a subset of the collected 1620 completions, specifically those in which both judges agreed that the subject of the completion unambiguously referred to either the Subject or the Object. We exclude from the analysis those completions in which the judges disagreed about the coding, or the completions coded as Join, Other and Unsure. Also discarded are the completions in which a Null was not used in the Null condition (1.3% (n=23) of the data). I take this low percentage as evidence that participants understood what a Null is. In total, 1098 completions were analyzed. These completions were further coded according to (i) the rhetorical relation between the two sentences and (ii) what type of referring expression was used in subject position in the Free condition: Null, Overt, proper name, etc. In this case, any disagreement was individually discussed and the judges agreed on a decision.

(11) summarizes the typology of rhetorical relations assumed in the paper (adapted from Kehler and Rohde (2013)):

- (11) a. Explanation: Infer that the second sentence describes a cause for the eventuality described in the first sentence.
“John handed a book to Bob. He no longer had a use for it.”
- b. Elaboration: Infer that both sentences provide descriptions of the same eventuality.
“John handed a book to Bob. He did so slowly and carefully.”

- c. Occasion: Infer a change of state from the second sentence, taking its initial state to be the final state of the eventuality described in the first sentence.
“John handed a book to Bob. He began reading it.”
- d. Result: Infer that the first sentence describes a cause or reason for the eventuality described in the second sentence.
“John handed a book to Bob. He thanked him for the gift.”
- e. Violated Expectation: Infer that the second sentence describes an unexpected result of the eventuality described in the first sentence.
“John handed a book to Bob. He showed no interest in reading it.”
- f. Parallel: Infer that the first and second sentences express similar eventualities, as if each provides a partial answer to a common question.
“John handed a book to Bob. He gave a magazine to him as well.”

To test for the statistical significance of the results, mixed-effect logistic regressions were performed, using R (R Core Team, 2013) and *lme4* (Bates et al., 2015). All models contained items and participants as random effects. Unless otherwise noted, the models also contained random slopes for all predictors and interactions (Barr et al., 2013). In case of non-converging models, the random effects structure was simplified, as specified for each individual model. Likelihood ratio tests are used to compare mixed-effects models different only in the presence or absence of the fixed effect in question. In the models with Prompt, which is a 3-level predictor, as a fixed effect, Free is treated as the baseline. In addition, we use pairwise comparisons and report the coefficient estimates (β estimates) and p-value for each combination.




Figure 1: Subject and Object continuations by Prompt

Figure 1 shows the percentage of subject and object reference in the three prompt conditions. In the three conditions we find an object bias, which is milder in the Null condition (62%), and greater in the Free (76%) and, particularly, Overt conditions (90%). To test for a main effect of prompt, a likelihood-ratio test was conducted between mixed-effects models different only in the presence or absence of a fixed main effect of prompt (see more details about the non-reduced model in Appendix 2, table 7). In both models, pronoun reference (subjects vs. object) was the dependent variable. The likelihood-ratio test showed a main effect of prompt ($\chi^2 = 105.99$, $p < .001$). Pairwise comparisons shows that the differences are significant in all three combinations: Free-Null ($\beta = -1.04$, $p < .001$),

Free-Overt ($\beta = 1.20$, $p = .006$) and Null-Overt ($\beta = 2.25$, $p < .001$). Thus, Hypothesis 1 is borne out: Nulls are more subject-biased than Overt, although overall both are object-biased in this context.

Figure 2 breaks down the data of the Free condition, showing the subject and object bias of those completions in which participants chose to use either a Null or an Overt. We can observe that the results in figures 1 and 2 are comparable: while Overt displays a strong bias towards the object (87%), Nulls show a much milder bias (57%). We compared two models in which pronoun reference (subject vs. object) was the dependent variable and which differed only in the presence or absence of a fixed main effect of pronoun type. A likelihood-ratio test showed the effect of pronoun type is significant ($\chi^2 = 23.90$, $p < .001$).




Figure 2: Subject and Object continuations by pronoun type in the Free condition




Figure 3: Distribution of referring expressions in the Free condition

Let us now examine what kind of referring expressions participants chose in the Free condition to refer to the previous subject and object. The data is summarized in figure 3.⁷ For the subject, the preferred form is the Null (72%) followed by a proper name (22%). For the object, the reverse order is found (30% of Nulls and 56% of proper names). The Overt pronoun is only the third most used form for both cases, accounting for only between 7% and 14% of the data. To test for

7. We exclude from the analysis the cases in which other forms, such as noun phrases or demonstratives were used, which accounted for less than 5% of the data.

the statistical significance of the data, we grouped the data in two binary categories: (i) choice of Null vs. not-Null as a referring expression and (ii) choice of name vs. not-name as a referring expression. We compared two models in which referring expression (Null vs. not-Null) was the dependent variable and which differed only in the presence or absence of a fixed main effect of pronoun reference (subject vs. object) (see more details about the non-reduced model in Appendix 2, table 8). A likelihood-ratio test showed the effect of pronoun reference is significant ($\chi^2 = 63.44$, $p < .001$). We further compared two models in which referring expression (name vs. not-name) was the dependent variable and which differed only in the presence or absence of a fixed main effect of pronoun reference (subject vs. object) (see more details about the non-reduced model in Appendix 2, table 9). A likelihood-ratio test showed that the effect of pronoun reference is significant ($\chi^2 = 63.63$, $p < .001$).

The asymmetry between interpretation and production observed by English has been replicated in the Catalan data for Nulls. Although Nulls displayed a mild interpretation bias towards the object, they display a very strong production bias towards the subject. When participants could choose a form to refer to the subject, they overwhelmingly chose a null pronoun. However, when they had to interpret a pronoun, the pragmatic biases came into play and overwhelmed the production bias. Given that it was more likely that the object (and not the subject) was mentioned next, the pronoun was interpreted with an object bias. In the next section, I discuss how this data fits in the Bayesian model proposed by Rohde and Kehler (2014).

The data has also uncovered another asymmetry concerning Overts. Although interpreters have a very strong bias to interpret the overt pronoun as referring to the object, it is clearly not the preferred form to refer to the object. More discussion of this asymmetry is postponed to section 3.4.

Let us now examine the effect of rhetorical relations. First, table 1 shows the distribution of rhetorical relations in each condition. We can take the results of the Free condition as the neutral results for TPV contexts: that is, as an estimate of which kind of rhetorical relations we are mostly likely to encounter after a TPV context. The data shows that, in a TPV context, we can expect a fair share of Occasions and Explanations, followed by Results, Elaborations and Violated Expectations.⁸ Now, observe how forcing participants to use either a Null or an Overt causes a switch in the percentage of observed rhetorical relations. Null prompts cause a rise of the rhetorical relations with subject bias, Elaboration and Explanation, and a decrease of the rhetorical relations with object bias, Occasion and Result. The opposite result is found for Overts.

	Free	Null	Overt
Explanation	25	33	17
Elaboration	18	32	9
Occasion	27	15	46
Result	15	10	20
Violated Expectation	12	9	6
Parallelism	3	1	2
Total	100	100	100

Table 1: % rhetorical relations by Prompt.

8. Since Parallelism occurred very rarely, it will not be discussed further.




Figure 4: Distribution of rhetorical relations by Prompt

In order to see the pattern more clearly, rhetorical relations can be grouped in two categories: subject-biased relations, which include Elaboration and Explanation, and object-biased relations, which include Result and Occasion. Figure 4 shows the distribution of these two categories by prompt. While in the Free condition, subject and object-biased relations are evenly split, Nulls clearly favor subject-biased relations and Overts clearly favor object-biased relations.

In order to test for statistical significance, two models were compared in which rhetorical relation bias (subject-biased vs. object-biased, as described above) was the dependent variable and which differed only in the presence or absence of a fixed main effect of prompt type (Null, Overt, Free) (see more details about the non-reduced model in Appendix 2, table 10). A likelihood-ratio test showed the effect of prompt type is significant ($\chi^2 = 142.92$, $p < .001$). Pairwise comparisons shows that the differences are significant in all three combinations: Free-Null ($\beta = 0.95$, $p < .001$), Free-Overt ($\beta = -1.30$, $p < .001$) and Null-Overt ($\beta = -2.25$, $p < .001$). Thus, hypothesis 2 is also borne out: type of prompt affects the distribution of rhetorical relations.




Figure 5: Subject continuations by rhetorical relation bias and pronoun type

Let us now turn to the question of whether pronoun interpretation is affected by rhetorical relations. Figure 5 shows the proportion of continuation about the subject by rhetorical relation bias in Conditions 1 and 2 (with Nulls and Overts). It can again be seen how rhetorical relations greatly affect the percentage of subject bias and the split between subject and object-biased relationships.

Object-biased relations show almost no reference to the subject, regardless of whether we find a Null or an Overt. In contrast, Subject-biased relations show a greater proportion of continuations about the subject.

Two models were compared in which pronoun reference was the dependent variable and which differed only in the presence or absence of a fixed main effect of rhetorical relation bias (subject-biased vs. object-biased) (see more details about the non-reduced model in Appendix 2, table 11). A likelihood-ratio test showed the effect of rhetorical relation is significant ($\chi^2 = 350.12$, $p < .001$). Hypothesis 3 is borne out: pronoun interpretation is affected by rhetorical relations.

While the results of object-biased relations is very homogenous, it is worth taking a closer look at the data for subject-biased relations. Table 2 summarizes the proportion of continuation about the subject in Explanation and Elaboration by pronoun type. While Elaborations shows a high percentage of subject references, the behavior of Explanation is somewhat unexpected. Although the proportion of subject continuations is higher than for object-biased relations, it is not clearly subject-biased either. This unexpected result will be discussed in section 3.4.

	Null	Overt
Explanation	29	32
Elaboration	72	50

Table 2: Proportion of continuations about the subject in Explanations and Elaborations by pronoun type

The results so far have confirmed that pragmatic factors affect pronoun interpretation. That is, pronouns biases are different across different rhetorical relations. For instance, Nulls are mostly interpreted as referring to the subject in cases of Elaborations, while they are interpreted as referring to the object in cases of Occasions. Now, we should take a look at the production data. Remember that the hypothesis is that the pronouns biases for a particular referent should be similar across different rhetorical relations. Figure 6 shows the percentage of Nulls referring to the subject in three rhetorical relations.⁹

We compared two models in which referring expression (Null Pronoun vs. No Pronoun) was the dependent variable and which differed only in the presence or absence of a fixed main effect of rhetorical relation (Elaboration being the baseline; see more details about the non-reduced model in Appendix 2, table 12). A likelihood-ratio test showed the effect of rhetorical relation is not significant ($\chi^2 = 0.25$, $p = .88$). Thus, hypothesis 4 is borne out: the production of Nulls remains constant and it is not affected by whether the rhetorical relation is subject biased or not.

3.3 Applying the Bayesian model

Let us now examine the predictions of the Bayesian model, presented in Section 2.3. In particular, we will examine the production and interpretation of the null pronoun to refer to a subject antecedent.¹⁰ We can, thus, rewrite the equation in (9) as (12).

9. We only present the cases in which there were at least 10 data points per rhetorical relation, which was not the case for Result, Violated Expectation and Parallelism.

10. An analysis for the overt pronoun will not be presented. There are not enough data points for a reliable estimation, given that it was scarcely produced in the Free conditions




Figure 6: Proportion of nulls referring to the subject by rhetorical relation

(12)

$$P(\text{subject} \mid \text{null}) = \frac{P(\text{null} \mid \text{subject}) P(\text{subject})}{P(\text{null})}$$

All these parameters can be estimated using data from the discourse-completion study. First, $P(\text{subject} \mid \text{null})$ is the probability that, given that a null pronoun has been used, it will be interpreted as referring to the subject. This probability, which can be estimated using the data of the Null condition, is 0.38 (see figure 1). Second, the reverse probability, $P(\text{null} \mid \text{subject})$, is the probability that a null pronoun will be used given that the referent is the subject. This probability, which can be estimated using data from the Free condition, is 0.72 (see figure 3). Third, the overall probability that the subject is the referent (regardless of the form used) is 0.24 (estimated using data from the Free condition, see figure 1). Finally, the denominator is the probability that a null pronoun is used, which can be calculated by summing the terms in the numerator by all possible referents (in this case subject and object). All the parameters are summarized in (13). The predicted probability is 0.44, not far from the observed probability (0.38). In order to test whether the correlation between observed and predicted probabilities is statistically significant we carried out a linear regression test over item means (that is, we computed the mean observed and predicted probabilities for each item).¹¹ The correlation is significant (adjusted $R^2 = 0.40$, $p = .007$).

- (13)
- a. Observed $P(\text{subject} \mid \text{null}) = 0.38$
 - b. $P(\text{null} \mid \text{subject}) = 0.72$
 - c. $P(\text{subject}) = 0.24$
 - d. $P(\text{null}) = P(\text{null} \mid \text{subject}) * P(\text{subject}) + P(\text{null} \mid \text{object}) * P(\text{object}) = 0.40$
 - e. Predicted $P(\text{subject} \mid \text{null}) = 0.43$

The Bayesian account outperforms other models, such as the Mirror Model, which equates $P(\text{subject} \mid \text{null})$ with the probability that a null is used given that the referent is the subject ($P(\text{subject} \mid \text{null})$), or the Expectancy model, which equates $P(\text{subject} \mid \text{null})$ with the probability that the subject is named again ($P(\text{subject})$). In the case of the Mirror Model the correlation is not significant (ad-

11. We excluded items for which in the Free condition participants did not refer to one of the two referents. That was the case for 1 of the 18 items.

justed $R^2 = -0.06$, $p = .71$); in the case of the Expectancy Model the correlation is significant (adjusted $R^2 = 0.37$, $p = .009$), but slightly lower than the one obtained with the Bayesian model.

3.4 Discussion

Experiment 1 mostly confirms the model by Kehler and Rohde (2014) about the role of rhetorical relations. Different pronoun prompts raise or lower the probability of particular rhetorical relations, which in turn carry their own interpretation biases. For instance, Occasion and Result are strongly object-biased and Elaboration is subject-biased. In contrast, pronoun production is not affected by such pragmatic factors: the rate of pronominalization is similar in Occasion and Elaboration, although they display opposite interpretation biases.

One of the results that does not fit precisely with their account is the low percentage of subject references with a null pronoun in the case of Explanation, which was only 29%. If Explanation is subject-biased and Nulls are subject-biased, why do 71% of Nulls in Explanation refer to objects? I believe the answer is that two different types of explanations are in play for subject and object references. (14) shows some examples of subject-referring explanations, while (15) shows some examples of object-referring explanations.

- (14) a. L'Elena li va regalar un llibre a la Mercè. \emptyset Sap que li encanten
The Elena DAT gave a book to the Mercè. \emptyset Knows that DAT please
aquesta mena de regals.
this kind of gifts.
'Elena gave a book to Mercè. (She) knows she adores this kind of gift.'
- b. El Joan va donar una joguina a l'Enric. \emptyset va preferir cedir abans que
The Joan gave a toy to the Enric. \emptyset preferred to give in before that
barallar-se un altre cop.
fight one other time
'Joan gave a toy to Enric. (He) preferred to give in rather than getting into a fight again.'
- (15) a. El Pere li va passar la clau anglesa al Marc. \emptyset La necessitava per obrir un
The Pere DAT passed the key English to the Marc. \emptyset it needed to open a
calaix enorme.
drawer huge.
'Pere passed the wrench to Marc. (He) needed it to open a huge drawer.'
- b. L'Alba li va deixar el cotxe a la Maria. \emptyset Havia d'anar als Pirineus a
The Alba DAT lent the car to the Maria. \emptyset had to go to the Pyrenees to
veure la seva família.
see the her family.
'Alba lent her car to Maria. (She) had to go to the Pyrenees to see her family.'

The object-referring completions answer a question like 'Why did the object referent **need** the transferred possession?' and typically include transferred possessions which are used to achieve something: a wrench is used to do something; a car is used to go somewhere. In a more fine-

grained typology of rhetorical relations, such examples could be coded as conveying purpose, rather than explanation.¹² In contrast, the subject-referring completions usually involved sentences with transferred possessions such as ‘book’ and ‘toy’. In this case, since it is not very informative to write about why the object referent needs a book or a toy, the completions rather explained why the subject transferred the possession to the object.

In fact, according to Kehler et al. (2008), the biases in the relations that express Cause or Effect (Explanation and Result, respectively) “will depend on the semantics incorporated in the passage and the referent to which causality or consequentiality is most likely to be attributed in a particular context” (page 26). Thus, Explanations are not intrinsically subject-biased: in some cases it is more likely to form a causal link between the two sentences with the subject antecedent, and in other cases with the object antecedent (see Bott and Solstad (2014) for a typology of explanations which attempts to make this idea more precise).

The data has also uncovered a strong asymmetry between the production and interpretation of the overt pronoun. Upon hearing an overt pronoun, a hearer will readily interpret it as referring to the object, even though the probability of using an overt pronoun to refer to an object is very low. This data shows that that the division of labor between Nulls and Overts is only partial: present in interpretation, but not in production. We, thus, see again that production and interpretation are not mirror images of each other.

The fact that Overts come only in third place in terms of the referring expressions chosen, after proper names and Nulls, suggests that, although they do display a clear object bias, their role in discourse is not to indicate such a bias, but rather to signal other relevant properties. Contrast seems the obvious candidate. As mentioned in Section 2, Overts often convey a contrastive flavor and are compulsory when they encode Focus or Contrastive Topics.

In order to examine the function of Overts more closely, we can examine their occurrence in the Free conditions; that is, those completions in which participants freely choose to use a pronoun. There were 38 such occurrences, of which 87% had object reference and 13% subject reference. In those cases, we do indeed find some examples in which Overts are conveying a Contrastive Topic, see (16), or a Focus, see (17), in which the Overt is in the postverbal Contrastive Focus position.

- (16) L'Elena li va regalar un llibre a la Mercè. Ella en canvi li va regalar una rosa.
The Elena DAT gave a book to the Mercè. She in change DAT gave a rose.

“Elena gave a book to Mercè. She gave him a rose instead.”

- (17) El Sergi li va facilitar totes les eines a l'Àlex. Així no havia de buscar-les ell.
The Sergi DAT supplied all the tools to the Àlex. Thus not had to search them he.
“Sergi supplied all the tools to Àlex. Thus, he was not the one that had to look for them”.

These cases, albeit possible, are not by any means representative of most of the data: that is, it is not the case that pronouns were mostly used to convey Contrastive Topic or Focus. There is, however, a wider notion of contrast that may be useful to understand the data. Table 3 shows the distribution of rhetorical conditions among the occurrences of Overts in the Free condition.

If we compare this distribution to the one of the whole Free condition (see table 1), we can see that the most striking difference is the higher percentage of Violated Expectations. In a Violated Expectation there is some kind of contrast between what is expected to happen and what really

12. I thank an anonymous reviewer for this observation.

Explanation	29
Occasion	29
Violated Expectation	21
Result	13
Parallelism	8

Table 3: Distribution of rhetorical conditions for Overts in the Free condition

happened, as illustrated in (18) with some of the completions of the experiment. So, although these cases are not cases of Foci or Contrastive Topics, they do convey contrast at a discourse level. The use of Overts is, thus, favored by discourse contrastivity, at least in this context, although this is not a necessary condition for their appearance. The production of Overts will be further discussed in connection to Experiment 2, which is presented in the next section.

- (18) a. La Gemma li va subministrar tot el material necessary a l’Elisena. Tot i
The Gemma DAT supplied all the material necessary to the Elisenda. All and
això, ella no li va agrair.
this, she not DAT thanked.
“Gemma supplied all the necessary material to Elisenda. However, she did not thank her.
- b. La Sra. Molins li va cedir la seva col·lecció de segells a la Lluïsa. Però ella
The Ms. Molins DAT bestowed the her collection of stamps to the Lluïsa. But she
no sabia què fer-ne.
not know what do PART.
“Ms. Molins bestowed her stamp collection on Lluïsa. But she did not know what to do with it.

4. Experiment 2: implicit causality verbs

Experiment 2 is also a discourse-completion study, which uses implicit causality verbs (ICVs). As mentioned, ICVs are useful to study the role of semantico-pragmatic factors, since they attribute the cause of the event they denote either to the subject (ICV1s) or to the object (ICV2s), which has been shown to affect how the pronouns that follow these verbs are interpreted.

4.1 Methods

Materials

Experiment 2 is similar to experiment 1, but it contains two factors: (i) verb type (ICV1 and ICV2) and (ii) prompt type (Null, Overt, Free Prompt (Free Condition)). There were, thus, 6 conditions: (19) illustrates the three conditions with an ICV1 and (20) the three conditions with an ICV2.

- (19) a. Condition 1: ICV1 + Null
La Núria va sorprendre la Maria. Ø ...
‘Núria surprised Maria. Ø ...’

- b. Condition 2: ICV1 + Overt
La Núria va sorprendre la Maria. Ella ...
'Núria surprised Maria. She ...'
 - c. Condition 3: ICV1 + Free
La Núria va sorprendre la Maria. ...
'Núria surprised Maria. ...'
- (20)
- a. Condition 4: ICV2 + Null
La Núria va felicitar la Maria. \emptyset ...
'Núria congratulated Maria. \emptyset ...'
 - b. Condition 5: ICV2 + Overt
La Núria va felicitar la Maria. Ella ...
'Núria congratulated Maria. She ...'
 - c. Condition 6: ICV2 + Free
La Núria va felicitar la Maria. ...
'Núria congratulated Maria. ...'

The experiment contained 30 items: 15 with an ICV1 and 15 with an ICV2. In each of the items, two referents of the same gender were mentioned, half containing masculine names and the other half feminine names. The first referent always appeared as subject and the second one as either direct object, indirect object or predicative complement. From now on, I will be referring the non-subject argument as 'object'. Three lists were constructed, so that each participant only saw each item with one of the prompts. The lists also contained 20 fillers with non-ICV verbs and with a connector or a temporal expression as a prompt. The full list of critical items with their English translations can be seen in Appendix 1.

Based on the previous literature, we formulate the following hypotheses:

1. Nulls are expected to receive more subject interpretations than Overt.
2. Given their lexical semantics, ICVs are expected to trigger a high number of Explanation completions.
3. Pronoun interpretation is expected to be affected by pragmatic factors. For Explanations, we expect a contrast between ICV1s and ICV2s: ICV1s are predicted to be subject-biased and ICV2s to be object-biased. We do not expect other rhetorical relations to be sensitive to the contrast between ICV1 and ICV2. As we found in Experiment 1, Elaboration is expected to be subject-biased, and Occasion and Result are expected to be object-biased, regardless of verb type.
4. Pronoun production is not expected to be affected by rhetorical relations. In the Free condition, the rate of pronominalization for a particular referent is expected to be similar in ICV1s and ICV2s.

Procedure

The same procedure described in the previous experiment was followed.

Participants

Seventy-eight participants took part in the experiment. None of the participants had participated

in Experiment 1. They were all native speakers of Catalan and students at the Universitat Pompeu Fabra and were entered in a raffle to win a gift certificate.

4.2 Results

A total of 2340 completions (30 items * 78 participants) were collected. The data was coded following the same procedure explained for Experiment 1. The antecedent of the referent first mentioned in the completion was coded into one of the following categories: Subject, Object, Joint (plural reference to both the subject and the object), Other or Unsure (if the pronoun could be understood as referring to more than one referent). The results are based on a subset of the completions, in which both judges agreed that the subject unambiguously referred to the previous subject or the previous object. The same exclusions reported for the previous experiment were carried out: all those completions in which the judges did not agree or the subject was coded as Unsure, Joint or Other. Also discarded were the cases in which a null pronoun was not used in the conditions ICV1+Null and ICV2+Null, which amounted to 0.68% (n=16) of the data. In total, 1934 completions were analyzed. These completions were coded according to the rhetorical relation between the two sentences and type of referring expression (Null, Overt, Proper Name, etc.) used in the two Free conditions. As in Experiment 1, any disagreement between the judges was individually discussed and the judges agreed on a decision. The statistical analysis followed the same methods previously described for Experiment 1.

Let us start by discussing hypothesis 1, which predicts that Nulls should be more subject-biased than Overts. Figure 7 shows that this is indeed the case both for ICV1s and ICV2s. We compared two models in which pronoun reference (subject vs. object) was the dependent variable and which differed only in the presence or absence of a fixed main effect of pronoun type (Null vs. Overt). A likelihood-ratio test showed that the effect of pronoun type is significant ($\chi^2 = 175.84$, $p < .001$). If we take the data in the Free conditions, a similar pattern emerges, as can be seen in figure 8. Again, in a likelihood-ratio test comparing two models which differ only in the presence or absence of a fixed main effect of pronoun type, the effect of pronoun type is significant ($\chi^2 = 21.78$, $p < .001$).




Figure 7: Proportion of continuations about the subject by verb type (ICV1 vs. ICV2) and prompt type (Null vs. Overt)

The two figures above have also uncovered another interesting pattern: there are more subject references with ICV1s than with ICV2s, both with Overt and Nulls. As a result, Nulls show a clear bias towards the subject with ICV1s, and Overts a strong bias towards the object with ICV2s. In the other two combinations (Null + VC2, and Overt + VC1), the two biases conflict with each




Figure 8: Proportion of continuations about the subject in the Free condition by verb type and form type

other and, as a result, there is no clear tendency. A model with verb type and pronoun type as fixed effects was compared to a model with an interaction between pronoun type and verb type (to achieve convergence, both models only contained items and participants as random effects; see more details about the non-reduced model in Appendix 2, table 13). A likelihood-ratio test showed that the interaction between pronoun type and verb type is significant ($\chi^2 = 7.98$, $p < .005$).

We can understand these results as a by-product of both hypotheses 2 and 3: we are expecting ICV contexts to trigger many Explanations, and Explanation is the rhetorical relation which imposes different biases (ICV1s towards the subject and ICV2s towards the object). Thus the prevalence of Explanations is what is responsible for the pattern observed in the tables above. Recall that the prediction is that the difference in biases in ICV1 and ICV2 contexts should only be present in Explanation and, therefore, should disappear with other relations.

Let us start by seeing whether ICVs really triggered a high number of Explanation continuations (Hypothesis 2). Table 4 shows the distribution of rhetorical relations: Explanation does indeed dominate, followed by Elaboration and Result. Hypothesis 2 is indeed borne out.

Explanation	61
Elaboration	19
Result	14
Violated Expectation	3
Occasion	2
Parallelism	1

Table 4: Distribution of rhetorical relations

We can now examine whether pronoun interpretation is affected by pragmatic factors. Figures 9 and 10 show the subject bias of Nulls and Overts respectively by verb type and rhetorical relation.¹³

Two models were compared in which pronoun reference was the dependent variable. One of the models had rhetorical relation, verb type and pronoun type as fixed effects, while the other had an interaction between the three fixed effects (Elaboration being the baseline; to achieve convergence, both models only contained items, participants and pronoun type as random effects; see more details about the model with the interaction in Appendix 2, table 14). A likelihood-ratio test showed that the

13. We eliminate from the analysis the rhetorical relations which account for less than 5% of the data.




Figure 9: Proportion of continuations about the subject with Nulls (by rhetorical relation and verb type)




Figure 10: Proportion of continuations about the subject with Overts (by rhetorical relation and verb type)

interaction is significant ($\chi^2 = 175.54, p < .001$). Having seen that there is a significant interaction, let us examine the data in specific rhetorical relations. For each rhetorical relation, two models were compared in which pronoun reference was the dependent variable and which differed only in the presence or absence of a fixed main effect of verb type. This comparison was done both for the data in the Null and in the Overt condition. As expected, in both cases, likelihood-ratio test showed the effect of verb type was significant for Explanations: ICV1s are subject biased and ICV2s are not (in the Null condition, $\chi^2 = 30.80, p < .001$; in the Overt condition, $\chi^2 = 42.21, p < .001$). Elaboration is subject-biased in both types of contexts, both with Nulls and Overts: the effect of verb type is not significant for Overts ($\chi^2 = 2.24, p = .52$), but it is significant for Nulls ($\chi^2 = 25.63, p < .001$). Result yields an unexpected significant contrast between ICV1 and ICV2 for the Null pronoun condition, ($\chi^2 = 14.46, p < .001$) while no contrast arises with the Overt pronoun ($\chi^2 = 1.88, p = 0.59$). I will discuss this unexpected contrast in Section 4.4. With the exception of the behavior of Result with Nulls, hypothesis 3 is also borne out: pronoun interpretation is affected by pragmatic factors.

Finally, let us turn to production data by examining what referring expression participants used to refer to subject and object depending on whether the context sentence contained an ICV1 or an

ICV2.¹⁴ Figures 11 and 12 summarize the data. The main observation is that, in both graphs, the distribution of referring expressions is fairly similar. In figure 11 we can observe, that to refer to the subject, participants overwhelmingly used a null pronoun regardless of whether the context was ICV1 or ICV2. In figure 12 we can see that, to refer to the object, the preferred form was also the null pronoun, but the percentage of proper names increased as well. Whether the verb was ICV1 or ICV2 does not seem to make a difference. Again we find that in both cases Overt forms were scarcely produced.




Figure 11: Distribution of referring expressions to refer to the subject




Figure 12: Distribution of referring expressions to refer to the object

We compared two models in which referring expression (Null vs. not-Null) was the dependent variable. One of the models had verb type and reference as fixed effects, while the other had an interaction between the two fixed effects (to achieve convergence, both models only contained items, participants and reference as random effects; see more details about the non-reduced model in Appendix 2, table 15). A likelihood-ratio test showed the effect of the interaction is not significant ($\chi^2 = 0.25$, $p = .61$). In fact, as it can be seen in 15, only reference is significant, while neither verb type nor the interaction between verb type and reference are significant. Thus, hypothesis 4 is borne out: production is not affected by rhetorical relation, which greatly affects interpretation.

14. As in experiment 1, cases in which a noun phrase or a demonstrative was used are excluded from the analysis.

4.3 Applying the Bayesian model

Let us examine again how well the Bayesian model can predict the observed interpretation biases of the Null pronoun to refer to a subject antecedent.

The interpretation biases observed in the data and the ones predicted by the Bayesian models (computed as explained in 3.3 for experiment 1) are summarized in (21). It can be seen how the Bayesian model adequately captures the tendencies in the data. A linear model analysis was performed over item means, and the correlation between the expected and the observed probabilities is significant (adjusted $R^2=0.57$, $p < .001$). The Bayesian account outperforms the the Mirror Model, which equates $P(\textit{subject} | \textit{null})$ with the probability that a null is used given that the referent is the subject ($P(\textit{subject} | \textit{null})$), or the Expectancy model, which equates $P(\textit{subject} | \textit{null})$ with the probability that the subject is named again ($P(\textit{subject})$). In the case of the Mirror Model the correlation is not significant (adjusted $R^2= -0.02$, $p= .65$); in the case of the Expectancy Model the correlation is significant (adjusted $R^2= 0.44$, $p < .001$), but lower than the one obtained with the Bayesian model.

- (21)
- a. **Observed** $P(\textit{subject} | \textit{null}) = 0.63$
 - b. $P(\textit{null} | \textit{subject}) = 0.89$
 - c. $P(\textit{subject}) = 0.50$
 - d. $P(\textit{null}) = P(\textit{null} | \textit{subject}) * P(\textit{subject}) + P(\textit{null} | \textit{object}) * P(\textit{object}) = 0.70$
 - e. **Predicted** $P(\textit{subject} | \textit{null}) = 0.64$

4.4 Discussion

Experiment 2 clearly showed how the interpretation of Nulls and Overt is influenced both by grammatical and pragmatic factors. Null subject bias increases in Explanations with an ICV1 verb, while it decreases with an ICV2 verb. The opposite is true for Overt: they have a strong object bias in Explanations with an ICV2 verb, which decreases in an ICV1 context. In contrast, pronoun production is not affected by pragmatic factors. The rate of use of Nulls in Explanations remains constant regardless of whether the verb was ICV1 or ICV2.

A surprising result of Experiment 2 was the bias shown by Nulls in Result relations. We expected Results to not be sensitive to verb type and, considering what was found with Transfer of Possession Verbs, to display an object bias. Instead, we found an object bias with ICV1s, and a subject bias with ICV2s. This behavior is actually not so surprising if we take into account that many verbs also display biases attributing the consequences of the event to one of their arguments. This bias is usually called ‘implicit consequentiality’¹⁵ (Stewart et al., 1998; Crinean and Garnham, 2006; Pickering and Majid, 2007). Many of the ICVs used in the experiments are actually psychological verbs (47%). In those cases, the subject position of an ICV1-sentence is occupied by the stimulus and the object by the experiencer, as in ‘intimidate’ or ‘terrify’ (see the first sentences in (22)). This pattern is reversed in ICV2s, such as ‘hate’ or ‘fear’ (see the first sentences in (23)). When a Result relation was expressed, it most often conveyed the consequences for the experiencer: this amounts to object references for ICV1s and subject references for ICV2s. The second sentences in (22) and (23) show typical completions of both cases. Thus, our hypothesis that, in Results, we would find an object bias (like we did for TPVs) was not borne out: instead, what we find is that Results display an experiencer bias (see Crinean and Garnham (2006) for more discussion on the relationship between

15. I thank an anonymous reviewer for pointing out this concept to me.

implicit causality, implicit consequentality and thematic roles). Table 5 shows the subject bias in Results by the thematic role of the subject: if the subject was an experiencer, the pronoun had a categorical subject preference, while if it was a stimulus, it had a strong object preference. In the cases of non-psychological verbs, we find the expected object preference.

(22) ICV1

- a. El Pol intimida l'Àlex. Ø Sempre que el veu Ø marxa ràpid.
The Pol intimidates the Àlex. Ø always that him sees Ø leaves quickly.
“Pol intimidates Àlex. Whenever (he) sees him, he leaves quickly.”
- b. La Pilar té la Blanca atemorida. Ø Sempre arriba plorant a casa.
The Pilar has the Blanca terrified. Ø Always arrives crying at home.
“Pilar terrifies Blanca. (She) always comes home crying.”

(23) ICV2

- a. En Guillem té por del Nicolau. Per això, Ø no va a l'escola.
The Guillem has fear o the Nicolau. For this, Ø not goes to the school.
“Guillem is afraid of Nicolau. This is why, (he) does not go to school.”
- b. La Candela té enveja de la Júlia. Per això, Ø vol deixar-la en ridícul.
The Candela has envy of the Júlia. For this, Ø wants leave her in embarrassment.
“Candela is envious of Júlia. This is why, (she) wants to ridicule her.”
- c. El Julià odia el Nil. Ø M'ha conessat que un dia Ø el matarà.
The Julià hates the Nil. Ø DAT has confessed that one day Ø him kill.
“Julià hates Nil. (He) has confessed to me that one day (he) will kill him.”

Agent	35
Experiencer	100
Stimulus	6

Table 5: % of subject references in Results by thematic role of the subject

Experiment 2 also confirmed the low probability of using an Overt to refer to either antecedent. Again, although presumably speakers could use an Overt to signal they do not want to refer to the most prominent referent, they usually do not do that and use a proper name instead. In order to understand why an Overt is used, we can examine those completions in which an Overt was chosen in the Free conditions. There were 49 such cases: they mostly refer to the object (80%), consistent with what occurred in the Overt conditions, and mostly occur with ICV2s (63%), which is what we would expect considering their object bias.

The rhetorical relations found in those cases are shown in table 6. The main finding is that there is an increase in the percentage of Results and a decrease in the percentage of Elaborations. Although we do find a few examples of Violated Expectation (see the examples in (24)), they do not account for a significant amount of the data, unlike what we found in Experiment 1. The reason is probably that the items in Experiment 1 favored completions in which participants narrated the expected (Occasions) or unexpected (Violated Expectations) events that followed the transfer of possession. In contrast, the items in Experiment 2 favor mainly Explanations about the event of the first sentence, or Results for the experiencer/theme. In all the examples in which the Overt was

used to express a Result, it referred to the object, which was either the experiencer or the theme (see examples in (25)).

Explanation	58
Elaboration	10
Result	24
Violated Expectation	6
Parallelism	2

Table 6: Distribution of rhetorical relations Free + Overt

- (24) a. L'Andreu va demanar disculpes al Joan. Tot i això, ell no
 The Andreu gave apologies to the Joan. Everything and this, he not
 va acceptar-les.
 accepted them
 “Andreu apologized to Joan. However, he did not accept them.”
- b. L'Iris confia en la Sílvia. I ella la va traïr.
 The Iris trusts in the Sílvia. And she her betrayed.
 “Iris trusts in Sílvia. And she betrayed her.”
- (25) a. L'Esteve va espantar el Roger. Ell el va empenyer com a venjança.
 The Esteve scared the Roger. He him pushed as a revenge.
 “Esteve scared Roger. He pushed him in revenge.”
- b. L'Adam va elogiar en Mateu. Ell es va posar vermell.
 The Adam praised the Mateu. He REFL turned red.
 “Adam praised Mateu. He blushed.”

5. Conclusion

This paper has uncovered two main asymmetries concerning pronouns in Catalan: (i) the interpretation of Nulls is affected by several pragmatic factors, which do not influence production, and (ii) while Overt pronouns show a strong interpretation bias towards the object, they are not frequently produced with this goal. Thus, the data supports only partially the idea that there is a division of labor between Null and Overt pronouns: while Nulls are clearly the default pronouns in Catalan, the role of the Overts is severely restricted.

Overall the data supports the model put forward by Rohde and Kehler (2014), according to which interpretation and production are not affected by the same set of factors. In their words, although it might be natural to expect “that speakers will employ pronouns in just those contextual circumstances in which the intended referent will be favored by the comprehender’s own biases” (Rohde and Kehler, 2014, p. 924), this is not what our data shows. The results are compatible with a model in which production is affected by structural factors (i.e. Nulls are subject biased and Overts are non-subject biased) and insensitive to semantico-pragmatic factors.

According to Fukumura and Van Gompel (2010), this difference between structural and semantico-pragmatic factors arises because only the former contributes to an entity’s accessibility. This makes

sense if we consider that while speakers can manipulate the structure of a sentence (for instance, choosing the passive form instead of the active form) depending on the relative accessibility of the entities in their discourse model, no such choice occurs with semantico-pragmatic factors. The speaker will use an ICV1 or an ICV2 depending on the meaning she wants to communicate; she will not choose one verb or another depending on the accessibility of the entities. The same reasoning can be applied to rhetorical relations: a speaker will link two utterances in her discourse by an Occasion or by an Elaboration depending on what is relevant for the discourse. It is unlikely that she will choose to elaborate on a previous utterance just to create a bias towards a certain entity. Now, by uttering an Elaboration a bias is certainly created and the interpreter can use this cue, but the reason to utter an Elaboration (as opposed to, say, an Occasion) is not to create the bias.

More generally, the picture that seems to emerge is one in which the interpreter combines multiple cues to assign reference (in particular, the grammatical bias associated with the choice of referential form with the prior probability of who will be mentioned), while the speaker ignores some of the cues that could potentially shape production. The result that production ignores factors which do affect interpretation challenges the *audience design* hypothesis (Clark, 1996), the idea that speakers always plan their utterance with the hearer in mind. It is instead compatible with a view that speakers for the most part use their own model to plan their utterances, while ignoring the hearer's (see Fukumura and van Gompel (2012) for experimental evidence showing that speakers use their own discourse model, and not the hearer's, when producing referential expressions).

A possible, although at this stage speculative, explanation of why speakers should ignore some cues which are useful for hearers would be that production is a more costly process than interpretation. A speaker has the burden to plan and produce the utterance, which includes selecting the appropriate lexical items within a huge lexicon, giving them a grammatical structure and articulating the relevant sounds. Thus, language production is undoubtedly hard (see, for instance, MacDonald (2013) for a model in which linguistic form is derived from the attempts by the speakers to mitigate utterance planning difficulties). In contrast, the task of the hearer is relatively simpler and that could be why the burden of integrating pragmatic information is passed exclusively to him. This would also be compatible with the finding that interpreters accommodate the needs of the speaker (from perspective-taking (Duran et al., 2011), to phonetic adaptations (Kraljic et al., 2008) or lexical and syntactic ambiguity resolution (MacDonald et al., 1994)).

Apart from providing support to the Bayesian approach with data from a language other than English¹⁶, this study also contributes to the characterization of Overts in Romance null-subject languages. Although we replicated the non-subject interpretation bias often discussed in the literature, a very striking result of our experiments is that Overts are rarely produced to fulfill this goal. We explored the possibility that Overts, like stressed pronouns in English, are only used in strictly contrastive uses, but this idea was not supported by the data. In fact, although, as mentioned, Overts were not used often, in some of the cases in which they were chosen, their function seemed to be precisely to refer to low-salience referents (which is fully compatible with their interpretation bias). But this raises the question: if Overts can be used to refer to low-salience referents, why don't speakers do it more often? I do not have a full answer to this question, but the difference in distribution of rhetorical relations seems to suggest they play an important role in licensing Overts. For future work, corpus-based research is planned so that the behavior of Overt pronouns can be further studied and the predictions of the Bayesian hypothesis can be tested with naturally-occurring

16. Similar studies were conducted in Japanese (Ueno and Kehler, 2016), but the predictions of the Bayesian approach were not tested.

data. A second venue for future work includes the use of richer contexts in the experiments. While our results support those who argue that predictability does not affect form production (Fukumura and Van Gompel, 2010; Rohde and Kehler, 2014), recent research has made the opposite point. In particular, in Rosa and Arnold (2017), the effects of predictability on form production were seen more clearly in tasks with very rich contexts, in which participants were asked to describe a series of pictures which told a coherent story. Thus, further work is needed to clarify the relationship between predictability and form production. Finally, we are also planning to further examine TPVs, comparing those in which the source argument is in subject position with those in which it is not, to corroborate that semantico-pragmatic factors do not play a role in determining referential form. We expect that we should not find a difference in pronominalization rate depending on whether the subject is the source or the goal, in the same way the we found similar pronominalization rates for rhetorical relations with opposed biases.

Acknowledgements

I am grateful to the three anonymous reviewers and to the Associate Editor for their constructive and detailed comments, suggestions and criticisms, which greatly contributed to improve this paper, as well as to the audiences of the venues where previous versions of this research was presented: Glif, Amore, Universitat de les Illes Balears, University of Vienna and XPrag 2017. I would also like to thank Aina Obis for her work in annotating the data, Andrew Kehler for many suggestions during the early stages of this project and Hannah Rohde for help in generating the graphs. The research underlying this article has been partially supported by projects FFI2015-66732-P and FFI2015-67991-P, funded by the Ministry of Economy, Industry and Competitiveness and the European Regional Development Fund (FEDER, UE).

Appendix 1

Experiment 1

El Joan li va portar un got d'aigua al Robert. ('Joan brought a glass of water to Robert.')

El Pere li va passar la clau anglesa al Marc. ('Pere passed a wrench to Marc.')

En Roger li va lliurar el treball al professor Ramos. ('Roger submitted his work to professor Ramos.')

En Pep li va enviar el seu CV al Toni. ('Pep sent his CV to Toni.')

La Marta li va donar una samarreta a la Ruth. ('Marta gave a t-shirt to Ruth.')

L'Elena li va regalar un llibre a la Mercè. ('Elena gifted a book to Mercè.')

La Jèssica li va servir l'arròs a la Carme. ('Jèssica served the rice to Carme.')

L'Eva li va tornar la grapadora a l'Esther. ('Eva returned the stapler to Esther.')

El Martí li va proporcionar medicines al Miquel. ('Martí provided medicines to Miquel.')

La Rosa li va vendre una postal a la Dolors. ('Rosa sold a postcard to Dolors.')

El Jordi li va dur el sopar a l'Ernest. ('Jordi brought dinner to Ernest.')

La Sra. Molins li ha cedit la seva col·lecció de segells a la Lluïsa. ('Mrs. Molins brought her stamp collection to Lluïsa.')

L'Adrià li va entregar un sobre a l'Albert. ('Adrià delivered an envelope to Albert.')

L'Alba li ha deixat el cotxe a la Marina. ('Alba lent her car to Marina.')

La Teresa li va acostar la sal a la Núria. ('Teresa moved the salt closer to Núria.')

El Sergi li va facilitar totes les eines a l'Àlex. ('Sergi provided all the tools to Àlex.')

La Gemma Martínez li va subministrar el material necessari a l'Elisenda. ('Gemma Martínez supplied the necessary material to Elisenda.')

El Joan li va donar una joguina a l'Enric. ('Joan gave a toy to Enric.')

Experiment 2

ICV1:

L'Andreu va demanar disculpes al Joan. ('Andreu apologized to Joan.')

La Mar va ofendre l'Irene. ('Mar offended Irene.')

La Raquel va enganyar l'Aurora. ('Raquel deceived Aurora.')

L'Àngels va humiliar la Marga. ('Àngels humiliated Marga.')

L'Abel va fer enfadar el Cesc. ('Abel annoyed Cesc.')

La Sara fa riure molt la Laia. ('Sara amuses Laia.')

El Miquel treu el Gerard de polleguera. ('Miquel bothers Gerard.')

L'Amanda va deixar la Montse bocabadada. ('Amanda amazed Montse.')

L'Esteve fa posar nerviós el David. ('Esteve makes David nervous.')

El Quim va decebre el Pep. ('Quim disappointed Pep.')

L'Esteve va espantar el Roger. ('Esteve scared Roger.')

En Lluís va sorprendre el Víctor. ('Lluís surprised Víctor.')

El Jordi té el Pau absolutament captivat. ('Jordi captivated Pau.')

En Pol intimida l'Àlex. ('Pol intimidates Àlex.')

La Pilar té la Blanca atemorida. ('Pilar frightens Blanca.')

ICV2:

La Marina va consolar la Susanna. ('Marina comforted Susanna.')

La Isabel va felicitar la Júlia. ('Isabel congratulated Júlia.')

La Noemí va ajudar la Clara. ('Noemí helped Clara.')

El Dídac es va burlar del Xavier. ('Dídac mocked Xavier.')

L'Àlícia va calmar la Xènia. ('Àlícia calmed Xènia.')

L'Àdam va elogiar en Mateu. ('Àdam praised Mateu.')

La Cati va esbrincar la Conxita. ('Cati told Conxita off.')

En Tomàs va renyar el Josep Antoni. ('Tomàs scolded Josep Antoni.')

L'Aleix va donar les gràcies al Felip. ('Aleix thanked Felip.')

L'Isaac va corregir el Joel. ('Issac corrected Joel.')

En Guillem té por del Nicolau. ('Guillem fears Nicolau.')

La Candela té enveja de la Joana. ('Candela is jealous of Joana.')

El Julià odia el Nil. ('Julià hates Nil.')

L'Iris confia en la Sílvia. ('Iris trusts Sílvia.')

La Laura valora la Xènia. ('Laura values Xènia.')

Appendix 2

Variable	Estimate	Error	z-value	p-value
Null	1.04	0.27	3.85	0.0001
Overt	-1.20	0.39	-3.05	0.002

Table 7: Experiment 1 Results. Pronoun reference (subjects, object) by prompt (Null, Overt, Free)
Model: $reference \sim prompt + (prompt|subject) + (prompt|item)$

Variable	Estimate	Error	z-value	p-value
Subject	-22.85	9.79	-2.33	< .01

Table 8: Experiment 1 Results. Form of referring expression (Null, not-Null) by reference (subject, object)

Model: $nullBinary \sim reference + (reference|subject) + (reference|item)$

Variable	Estimate	Error	z-value	p-value
Subject	30.07	6.87	4.37	< .001

Table 9: Experiment 1 Results. Form of referring expression (name, not-name) by reference (subject, object)

Model: $nameBinary \sim reference + (reference|subject) + (reference|item)$

Variable	Estimate	Error	z-value	p-value
Null	-0.95	0.24	-3.92	< .0001
Overt	1.30	0.25	5.18	< .0001

Table 10: Experiment 1 Results. Rhetorical relations (subject-biased, not subject-biased) by prompt (Free, Null, Overt)

Model: $rhetoricalBinary \sim prompt + (prompt|subject) + (prompt|item)$

Variable	Estimate	Error	z-value	p-value
Object-biased	-4.54	0.38	-11.886	< .0001

Table 11: Experiment 1 Results. Reference (subject, object) by rhetorical relation (subject-biased, not subject-biased)

Model: $reference \sim rhetoricalBinary + (rhetoricalBinary|subject) + (rhetoricalBinary|item)$

Variable	Estimate	Error	z-value	p-value
Explanation	0.97	2.82	0.34	0.73
Occasion	0.31	4.64	0.06	0.94

Table 12: Experiment 1 Results. Form of referring expression (Null, non-Null) by rhetorical relation (Elaboration, Explanation, Occasion)

Model: $nullBinary \sim rhetorical + (rhetorical|subject) + (rhetorical|item)$

Variable	Estimate	Error	z-value	p-value
Overt	-1.20	0.18	-6.6	< .0001
ICV2	-1.20	0.26	-4.46	< .0001
Overt*ICV2	-0.75	0.26	-2.82	0.004

Table 13: Experiment 2 Results. Reference (subject, object) by pronoun (Null, Overt) * verb type (ICV1, ICV2)

Model: $reference \sim pronoun * type + (I|subject) + (I|item)$

Variable	Estimate	Error	z-value	p-value
ICV2	1.49	0.63	2.33	0.01
Explanation	1.19	0.42	2.81	0.004
Result	-3.02	0.73	-4.08	< .0001
Overt	-1.10	0.45	-2.44	0.01
ICV2*Explanation	-4.55	0.66	-6.83	< .0001
ICV2*Result	1.12	0.99	1.13	0.25
ICV2*Overt	-0.84	0.75	-1.12	0.26
Explanation*Overt	0.25	0.55	0.45	0.64
Result*Overt	0.03	0.93	0.03	0.97
ICV2*Explanation*Overt	0.11	0.92	0.12	0.90
ICV2*Result*Overt	-2.30	1.36	-1.68	0.09

Table 14: Experiment 2 Results. Reference (subject, object) by pronoun (Null, Overt) * verb type (ICV1, ICV2) * rhetorical relation (Elaboration, Explanation, Result)

Model: $reference \sim pronoun * type * rhetorical + (pronoun|subject) + (pronoun|item)$

Variable	Estimate	Error	z-value	p-value
ICV2	-0.66	0.65	-1.02	0.30
Subject	-3.09	0.78	-3.92	< .0001
ICV2*Subject	0.52	1.04	0.50	0.61

Table 15: Experiment 2 Results. Form of referring expression (Null, non-Null) by reference (subject, object) * verb type (ICV1, ICV2)

Model: $formBinary \sim reference * type + (reference|subject) + (reference|item)$

References

- Alonso-Ovalle, L., Solera, S. F., Frazier, L. and Clifton, C. (2002), 'Null vs. overt pronouns and the topic-focus articulation in Spanish', *Journal of Italian Linguistics* **14:2**, 151–169.
- Ariel, M. (1990), *Accessing noun-phrase antecedents*, London: Routledge.
- Arnold, J. E. (1999), Reference form and discourse patterns., PhD 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.
- Barr, D. J., Levy, R., Scheepers, C. and Tily, H. J. (2013), 'Random effects structure for confirmatory hypothesis testing: Keep it maximal', *Journal of memory and language* **68**(3), 255–278.
- Bates, D., Mächler, M., Bolker, B. and Walker, S. (2015), 'Fitting linear mixed-effects models using lme4', *Journal of Statistical Software* **67**(1), 1–48.
- Bott, O. and Solstad, T. (2014), From verbs to discourse: A novel account of implicit causality, in C. F.-H. B. Hemforth, B. Mertins, ed., 'Psycholinguistic approaches to meaning and understanding across languages', Springer, pp. 213–251.
- Brown, R. and Fish, D. (1983), 'The psychological causality implicit in language', *Cognition* **14**(3), 237–273.
- Cameron, R. (1992), Pronominal and null subject variation in Spanish: constraints, dialects, and functional compensation, PhD thesis, University of Pennsylvania.
- Caramazza, A., Grober, E., Garvey, C. and Yates, J. (1977), 'Comprehension of anaphoric pronouns', *Journal of Verbal Learning and Verbal Behavior* **16**, 601–609.
- Carminati, M. N. (2002), The processing of Italian subject pronouns, PhD thesis, University of Massachusetts.
- Clark, H. H. (1996), *Using language*, Cambridge University Press: Cambridge.
- Crawley, R. A., Stevenson, R. J. and Kleinman, D. (1990), 'The use of heuristic strategies in the interpretation of pronouns', *Journal of Psycholinguistic Research* **19**(4), 245–264.
- Crinean, M. and Garnham, A. (2006), 'Implicit causality, implicit consequentiality and semantic roles', *Language and Cognitive Processes* **21**(5), 636–648.
- DiEugenio, B. (1998), Centering in Italian, in A. K. J. M. Walker and E. Prince, eds, 'Centering Theory in Discourse', Oxford University Press, pp. 114–137.
- Duran, N. D., Dale, R. and Kreuz, R. J. (2011), 'Listeners invest in an assumed other's perspective despite cognitive cost', *Cognition* **121**(1), 22–40.
- Fukumura, K. and Van Gompel, R. P. (2010), 'Choosing anaphoric expressions: Do people take into account likelihood of reference?', *Journal of Memory and Language* **62**(1), 52–66.

- Fukumura, K. and van Gompel, R. P. (2012), ‘Producing pronouns and definite noun phrases: Do speakers use the addressee’s discourse model?’, *Cognitive Science* **36**(7), 1289–1311.
- Givón, T. (1983), *Topic continuity in discourse*, John Benjamins Publishing Company.
- Gundel, J. K., Hedberg, N. and Zacharski, R. (1993), ‘Cognitive status and the form of referring expressions in discourse’, *Language* pp. 274–307.
- Keating, G. D., VanPatten, B. and Jegerski, J. (2016), ‘Online processing of subject pronouns in monolingual and heritage bilingual speakers of mexican spanish’, *Bilingualism: Language and Cognition* **19**(01), 36–49.
- Kehler, A. (2002), *Coherence, Reference and the Theory of Grammar*, Stanford, CA, CSLI Publications.
- Kehler, A., Kertz, L., Rohde, H. and Elman, J. L. (2008), ‘Coherence and coreference revisited’, *Journal of semantics* **25**(1), 1–44.
- Kehler, A. and Rohde, H. (2013), ‘A probabilistic reconciliation of coherence-driven and centering-driven theories of pronoun interpretation’, *Theoretical Linguistics* **39**(1-2), 1–37.
- Kraljic, T., Samuel, A. G. and Brennan, S. E. (2008), ‘First impressions and last resorts: How listeners adjust to speaker variability’, *Psychological science* **19**(4), 332–338.
- Luján, M. (1985), Binding properties of overt pronouns in null pronominal languages, in P. K. W. H. Eilforth and K. Peterson, eds, ‘Proceedings of the Chicago Linguistics Society’, Vol. 21, pp. 424–438.
- Luján, M. (1999), Expresión y omisión del pronombre personal, in V. Demonte and I. Bosque, eds, ‘Gramática descriptiva de la lengua española’, Espala-Calpe, Madrid, pp. 1275–1316.
- MacDonald, M. C. (2013), ‘How language production shapes language form and comprehension’, *Frontiers in psychology* **4**, 226.
- MacDonald, M. C., Pearlmutter, N. J. and Seidenberg, M. S. (1994), ‘The lexical nature of syntactic ambiguity resolution.’, *Psychological review* **101**(4), 676.
- Mayol, L. (2010a), ‘Contrastive pronouns in null-subject romance languages’, *Lingua* **120**(10), 2497–2514.
- Mayol, L. (2010b), ‘Refining salience and the position of antecedent hypothesis: a study of catalan pronouns’, *University of Pennsylvania Working Papers in Linguistics* **16**(1), 15.
- Mayol, L. and Clark, R. (2010), ‘Pronouns in catalan: Games of partial information and the use of linguistic resources’, *Journal of Pragmatics* **42**(3), 781–799.
- Pickering, M. J. and Majid, A. (2007), ‘What are implicit causality and consequentiality?’, *Language and Cognitive Processes* **22**(5), 780–788.
- R Core Team (2013), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: <http://www.R-project.org/>

- Rigau, G. (1986), Some remarks on the nature of strong pronouns in null-subject languages, in I. Bordelois, H. Contreras and K. Zagona, eds, 'Generative Studies in Spanish Syntax', Foris, Dordrecht.
- Rohde, H. and Kehler, A. (2014), 'Grammatical and information-structural influences on pronoun production', *Language, Cognition and Neuroscience* **29**(8), 912–927.
- Rosa, E. C. and Arnold, J. E. (2017), 'Predictability affects production: Thematic roles can affect reference form selection', *Journal of Memory and Language* **94**, 43–60.
- Samek-Lodovici, V. (1996), Constraints on Subjects: An Optimality Theoretic Analysis, Ph.d. thesis, Rutgers University.
- Silva-Corvalán, C. (1977), A discourse study of word order in the Spanish spoken by Mexican-Americans in West Los Angeles, Master's thesis, University of California, Los Angeles.
- Smyth, R. (1994), 'Grammatical determinants of ambiguous pronoun resolution', *Journal of Psycholinguistic Research* **23**(3), 197–229.
- Stevenson, R. J., Crawley, R. A. and Kleinman, D. (1994), 'Thematic roles, focus and the representation of events', *Language and Cognitive Processes* **9**(4), 519–548.
- Stewart, A. J., Pickering, M. J. and Sanford, A. J. (1998), Implicit consequentiality, in 'Proceedings of the Twentieth Annual Conference of the Cognitive Science Society', Lawrence Erlbaum Associates, NJ, pp. 1031–1036.
- Ueno, M. and Kehler, A. (2016), 'Grammatical and pragmatic factors in the interpretation of Japanese null and overt pronouns', *Linguistics* **54**(6), 1165–1221.
- Vallduví, E. (1992), *The informational component*, Garland, New York.