L1 attrition in Italian:

Pronominal reference and (pseudo-)relative parsing ambiguities

Alexander Allan Cairncross
Churchill College
Section of Theoretical and Applied Linguistics
University of Cambridge

This dissertation is submitted for the degree of Doctor of Philosophy
June 2023

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other university or similar institution. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other university or similar institution. This thesis does not exceed the prescribed word limit of 80,000 words.

Alexander Allan Cairncross
June 2023

Abstract

L1 attrition in Italian: Pronominal reference and (pseudo-)relative parsing ambiguities

Alexander Allan Cairncross

The present thesis presents a series of experiments investigating potential (syntactic) attrition effects in Italian beyond traditional interface structures. This is motivated by a critical review of the previous literature on first language (L1) attrition primarily focusing on the Interface Hypothesis (Sorace & Filiaci 2006). From this, we argue for two relevant patterns: (i) the often discussed relaxation of discourse constraints for specific forms (e.g., overt pronouns in null-subject languages) and (ii) a tendency, which we believe has so far gone largely unnoticed, for speakers to rely more heavily on other (interpretive) biases already present in their L1. To explore this second pattern and disentangle the role of interpretive/processing biases from the syntax-discourse interface, we compare pronominal subjects which are often taken to sit at the syntax-discourse interface with (pseudo-)relative clause attachment ambiguities. We argue that this is an interesting comparison as such attachment ambiguities involve two ambiguities which are both conditioned by principles of computational efficiency at the level of the parser and are not related to discourse. In a language like Italian, when encountering an embedded CP of the relevant type, we will assume that the parser must decide whether the string contains a relative clause (RC) or a pseudorelative (PR) which are structurally distinct but string identical (PR-First Hypothesis, Grillo & Costa 2014). Moreover, in the case of a RC reading with multiple possible attachment sites, the parser must additionally decide where to attach the embedded CP (Late Closure, Frazier 1978). Should structures at more external interfaces (e.g., the syntax-discourse interface) be more vulnerable under attrition, we would expect the two phenomena to be affected deferentially under attrition. However, should attrition lead to an increased reliance on certain biases regardless of interface status, we might expect the two phenomena to pattern together.

To compare the effects of attrition on the offline interpretive biases for the two phenomena, we carried out a sentence interpretation task with globally ambiguous items of both types in which we manipulated the pronominal forms and pseudorelative availability. To compare online processing biases, we conducted a self-paced reading task in which we additionally forced pronominal resolution and 'RC' attachment. Given some unexpected results from the control group for the self-paced reading task, we additionally present a sister eye-tracking-while-reading task as a validation (albeit with a different methodology) of the control baseline. For the experiments in the present thesis, all participants reported growing up monolingually in Italy. However, experimental groups differed from control groups in that the former had immigrated to a majority English-speaking country (e.g., the UK) during adulthood and had lived there for a minimum of 5 years. Recent re-exposure to the L1 was controlled for.

Results from the offline interpretation and online processing of both phenomena indicated that, where attested, attrition always surfaced an increased reliance on L1 biases. Strikingly this holds even for the prime syntax-discourse interface phenomenon, overt subject pronouns in a null-subject language like Italian. As such, we argue that future accounts for the attrition of syntactic phenomena should not focus on the specific cognitive demands of particular interface phenomena (e.g., pronominal forms and the integration of contextual information) but instead should seek a more parsimonious account couched in terms of increased reliance on L1 biases in ambiguity resolution more generally.

Acknowledgements

Thank you to Ianthi Tsimpli for your guidance throughout my studies. In particular, thank you for regularly challenging my arguments while also giving me space to develop my own ideas. Thank you to Margreet Vogelzang both for your advice and for always making time to talk through my sometimes never-ending questions on every topic imaginable. Thank you to Anna Cardinaletti and members of the BemboLab for hosting me as well as your invaluable help in making the in-person data collection possible. Thank you to Mariano Felice for your technical help when I ran into server problems. Without your aid, I would not have been able to finish the data collection for the self-paced reading experiments in chapters 5 and 6. Thank you to Pietro Baggio and Gianluca Porta for reading countless Italian sentences. Thank you to Sana Kidwai, Li Nguyen, Andromachi Tsoukala, and Rita Pataki for proofreading various drafts. Thank you to all of my participants without whom, none of the experiments would have been possible. Thank you to all the anonymous reviewers and conference/seminar audience members for your useful feedback on various parts of this work as it progressed. Finaly, thank you to Maciej Lemanowicz for your support throughout.

Table of contents

Declaration ii

Abstract iii

Acknowledgements iv

Table of contents viii

List of figures x

List of tables xii

List of abbreviations xiii

1 Introduction 1
- 1.1 Introduction 1
- 1.2 Originality 4
- 1.3 Thesis structure 5

2 General background and research questions 7
- 2.1 Introduction 7
- 2.2 Pronominal subjects in Italian (and related languages) 7
- 2.3 Attrition 12
  - 2.3.1 A working definition 12
  - 2.3.2 The Interface Hypothesis 13
  - 2.3.3 Previous work on attrition in null-subject languages 15
    - 2.3.3.1 Attrition and interpretive biases 15
    - 2.3.3.2 Attrition and (surface) overlap with the L2 26
    - 2.3.3.3 Attrition and production 30
    - 2.3.3.4 Attrition and online processing 32
    - 2.3.3.5 Attrition and re-exposure effects 37
    - 2.3.3.6 Summary of previous results 40
- 2.4 (Pseudo-)relative clause parsing ambiguities 40
  - 2.4.1 The PR account 42
  - 2.4.2 Previous (monolingual) work on PRs 43
  - 2.4.3 Attrition and attachment biases 47
- 2.5 Research questions 49
- 2.6 Chapter summary 51

3 Attrition and interpretive biases in globally ambiguous items 52
- 3.1 Introduction 52
- 3.2 Research questions 52
- 3.3 Method 52
  - 3.3.1 Participants 52
  - 3.3.2 Stimuli 53
    - 3.3.2.1 Pronominal items 54
    - 3.3.2.2 PR/RC items 55
  - 3.3.3 Procedure 55

3.3.4 Data cleaning and analysis 56
3.3.5 Hypotheses and predictions 57
3.4 Results 58
3.4.1 Pronominal items 58
3.4.2 PR/RC items 59
3.5 Discussion 60
3.5.1 Pronominal items 60
3.5.2 PR/RC items 61
3.6 Chapter summary 64

4 Construction and evaluation of the locally ambiguous stimuli 65
4.1 Introduction 65
4.2 Development of the experimental materials 65
4.2.1 PR/RC items 66
4.2.2 Pronominal items 67
4.2.3 Segmentation 68
4.2.4 Interim summary 69
4.3 Evaluation of the experimental items 69
4.3.1 Consultation with native linguists 69
4.3.2 Material for the norming study 70
4.3.3 Participants 71
4.3.4 Procedure 71
4.3.5 Results 72
4.3.6 Discussion and selection of the final items 73
4.4 Chapter summary 73

5 Attrition and the Position of Antecedent Hypothesis in locally ambiguous items 74
5.1 Introduction 74
5.2 Research questions 74
5.3 Experiment 1: Self-paced reading 74
5.3.1 Method 74
5.3.1.1 Participants 74
5.3.1.2 Stimuli 75
5.3.1.3 Procedure 76
5.3.1.4 Data cleaning and analysis 76
5.3.1.5 Hypotheses and predictions 77
5.3.2 Results 78
5.3.2.1 Reading times in window 5: Overt pronouns 78
5.3.2.2 Reading times in window 6: Null pronouns 79
5.3.2.3 Response accuracy for comprehension questions: Null and overt pronouns 81
5.3.3 Interim discussion 81
5.4 Experiment 2: Eye-tracking-while-reading 81
5.4.1 Method 81
5.4.1.1 Participants 81
5.4.1.2 Stimuli 82
5.4.1.3 Procedure 82
5.4.1.4 Data cleaning and analysis 83

5.4.1.5 Hypotheses and predictions 84
5.4.2 Results 84
5.4.2.1 Eye-movement measures in interest area 5: Overt pronouns 84
5.4.2.2 Eye-movement measures in interest area 6: Null pronouns 84
5.4.2.3 Response accuracy for comprehension questions: Null and overt pronouns 88
5.4.3 Interim discussion 88
5.5 Experiment 3: Self-paced reading again 89
5.5.1 Method 89
5.5.1.1 Participants 89
5.5.1.2 Stimuli 90
5.5.1.3 Procedure 91
5.5.1.4 Data cleaning and analysis 91
5.5.1.5 Hypotheses and predictions 91
5.5.2 Results 92
5.5.2.1 Reading times in window 5: Overt pronouns 92
5.5.2.2 Reading times in window 6: Null pronouns 95
5.5.2.3 Response accuracy for comprehension questions: Null and overt pronouns 95
5.6 Discussion 98
5.7 Chapter summary 102

6 Attrition and the PR-First Hypothesis in locally ambiguous items 103
6.1 Introduction 103
6.2 Research questions 103
6.3 Experiment 1: Self-paced reading 104
6.3.1 Method 104
6.3.1.1 Participants 104
6.3.1.2 Stimuli 104
6.3.1.3 Procedure 105
6.3.1.4 Data cleaning and analysis 105
6.3.1.5 Hypotheses and predictions 106
6.3.2 Results 107
6.3.2.1 Reading times in the critical window 107
6.3.2.2 Reading times in the post-critical window 107
6.3.2.3 Response accuracy for comprehension questions 109
6.3.3 Interim discussion 110
6.4 Experiment 2: Eye-tracking-while-reading 111
6.4.1 Method 111
6.4.1.1 Participants 111
6.4.1.2 Stimuli 111
6.4.1.3 Procedure 112
6.4.1.4 Data cleaning and analysis 112
6.4.1.5 Hypotheses and predictions 113
6.4.2 Results 113
6.4.2.1 Eye-tracking measures in critical interest area 113
6.4.2.2 Eye-tracking measures in the post-critical interest area 116
6.4.2.3 Response accuracy for comprehension questions 118
6.4.3 Interim discussion 119

6.5 Experiment 3: Self-paced reading again 121
6.5.1 Method 121
6.5.1.1 Participants 121
6.5.1.2 Stimuli 122
6.5.1.3 Procedure 122
6.5.1.4 Data cleaning and analysis 122
6.5.1.5 Hypotheses and predictions 123
6.5.2 Results 123
6.5.2.1 Reading times in the critical window 123
6.5.2.2 Reading times in the post-critical window 127
6.5.2.3 Response accuracy for comprehension questions 127
6.6 Discussion 129
6.7 Chapter summary 131

7 General discussion and conclusions 132
7.1 Introduction 132
7.2 Summary 132
7.2.1 Background and aims 132
7.2.2 Experimental results 133
7.2.2.1 Interpretive biases 133
7.2.2.2 Processing biases 133
7.3 Discussion 135
7.4 Conclusion 147

A Data accessibility 149

B Globally ambiguous stimuli 150

C Locally ambiguous stimuli 154

D Supplementary tables and figures 164

References 175

List of figures

2.1 Example of the picture selection task used in Tsimpli, Sorace, Heycock & Filiaci (2004) for the item in (17a). 15
2.2 Sketch of the [±TS] specification of pronominal forms in monolingual Italian, monolingual English, and bilingual Italian. 16
2.3 Example of the picture selection task used in Tsimpli et al. (2004) for the item in (19). 17
2.4 An example image for the corresponding sentence in (24). 20
2.5 Example of the images used for the picture selection task in Martín-Villena (2023) for the item in (26b). 23
2.6 Example images used in Kaltsa, Tsimpli & Rothman (2015) for the item in (34). 34
2.7 Example of the images used for the self-paced reading task in Martín-Villena (2023) for the item in (36a). 36
2.8 Summary of (relevant) universal PR/RC parse options. 43

3.1 Subject response rates for the pronominal items by group, pronominal form, and direction with 95% confidence intervals. 58
3.2 High attachment rates for the PR/RC items by group and predicate with 95% confidence intervals. 60

4.1 Global naturalness rating (on a scale from 1 'very unnatural' to '5 very natural') for norming items with 95% confidence intervals. 72

5.1 Global average reading times for the fifth window of items with overt pronouns with 95% confidence intervals. 78
5.2 Global average reading times for the sixth window of items with null pronouns with 95% confidence intervals. 79
5.3 Average reading times over trial order on the log millisecond scale for the sixth window of items with null pronouns with 95% confidence intervals. 80
5.4 Average response accuracy (by participant) for null and overt pronominal items with 95% confidence intervals. 80
5.5 Eye-tracking measures for overt pronominal items in the fifth region with 95% confidence intervals. 85
5.6 Eye-tracking measures for null pronominal items in the sixth region with 95% confidence intervals. 86
5.7 Average response accuracy (by participant) for null and overt pronominal items with 95% confidence intervals. 88
5.8 Global average reading times for the fifth window of items with overt pronouns for the control and experimental groups with 95% confidence intervals. 93
5.9 Average reading times by antecedent over trial order in log milliseconds for the fifth window of items with overt pronouns with 95% confidence intervals. 94
5.10 Average reading times by antecedent and group over trial order in log milliseconds for the fifth window of items with overt pronouns with 95% confidence intervals. 94
5.11 Global average reading times for the sixth window of items with null pronouns for the control and experimental groups with 95% confidence intervals. 96
5.12 Average response accuracy (by participant) for the control and experimental groups broken down by pronominal type with 95% confidence intervals. 97

6.1 Global average reading times for the critical window with 95% confidence intervals. 108
6.2 Global average reading times for the post-critical window with 95% confidence intervals. 108

6.3 Average response accuracy (by participant) with 95% confidence intervals. 109
6.4 Eye-tracking measures for the critical region with 95% confidence intervals. 114
6.5 Eye-tracking measures for the post-critical region with 95% confidence intervals. 117
6.6 Average response accuracy (by participant) with 95% confidence intervals. 118
6.7 Global average reading times for the critical window for the control and experimental groups with 95% confidence intervals. 124
6.8 Global average reading times for the post-critical window for the control and experimental groups with 95% confidence intervals. 125
6.9 Average response accuracy (by participant) for the control and experimental groups with 95% confidence intervals. 128

D.1 Responses to globally ambiguous matrix-first pronominal items in chapter 3 with the [+topic shift (TS)] readings separated into 'object' and 'other'. 164
D.2 Responses to globally ambiguous embedded-first pronominal items in chapter 3 with the [+TS] readings separated into 'object' and 'other'. 165

List of tables

2.1 Percentage of SpecIP responses by group and pronoun in both sessions adapted from Gargiulo & van de Weijer's (2020) Table 1. 39

3.1 Mean language use by participants as percentages of a typical day with standard deviations. 53
3.2 Model outputs for matrix- and embedded-first pronominal sentence interpretations. 59
3.3 Pairwise comparisons for matrix-first pronominal interpretations with Holm–Bonferroni corrected *p*-values. 59
3.4 Model output for PR/RC sentence interpretations. 59

4.1 Language use as percentages of a typical day for the participants of the norming study. 71
4.2 Model output for global norming ratings. 73
4.3 Naturalness ratings (on a scale from 1 'very unnatural' to 5 'very natural') for the final 32 items per type. 73

5.1 Mean language use by participant as percentages of a typical day with standard deviations. 75
5.2 Model output for the reading times of items with overt pronouns in the fifth window. 78
5.3 Model output for the reading times of items with null pronouns in the sixth window. 79
5.4 Model output for response accuracy for null and overt pronominal items. 81
5.5 Mean language use by participant as percentages of a typical day with standard deviations. 82
5.6 Average FP, GP, and TT for overt pronominal items in ms for the fifth region with 95% confidence intervals. 84
5.7 Model outputs for overt pronouns in the fifth region. 87
5.8 Average FP, GP, and TT for null pronominal items in ms for the sixth region with 95% confidence intervals. 87
5.9 Model outputs for null pronouns in the sixth region. 87
5.10 Model output for response accuracy for null and overt items. 89
5.11 Mean language use by participant as percentages of a typical day with standard deviations for the experimental and control groups. 90
5.12 Model output for the reading times of items with overt pronouns in the fifth window for the experimental and control groups. 92
5.13 Model output for the reading times of items with null pronouns in the sixth window for the experimental and control groups. 95
5.14 Model output for response accuracy for the experimental and control groups. 98

6.1 Mean language use by participant as percentages of a typical day with standard deviations. 104
6.2 Model output for the reading times of the critical window. 107
6.3 Model output for the reading times of the post-critical window. 109
6.4 Model output for response accuracy. 110
6.5 Mean language use by participant as percentages of a typical day with standard deviations, repeated from Table 5.5. 111
6.6 Average FP, GP, and TT for the critical region with 95% confidence intervals. 113
6.7 Model outputs for the critical region. 115
6.8 Average FP and TT for the postcritical region with 95% confidence intervals. 116
6.9 Model outputs for the post-critical region. 116
6.10 Model output for response accuracy. 118
6.11 Mean language use by participant as percentages of a typical day with standard deviations for the experimental and control groups. 122

6.12 Model output for the reading times of the critical window for the experimental and control groups. 126
6.13 Model output for the reading times of the post-critical window for the experimental and control groups. 126
6.14 Model output for the response accuracy for the experimental and control groups. 129

B.1 PR/RC items borrowed from Grillo & Costa's (2014) second experiment with added translations. 151
B.2 Pronominal items adapted from Tsimpli et al.'s (2004) sentence interpretation task with added translations. 153

C.1 Critical PR/RC items. 158
C.2 Critical pronominal items. 163

D.1 Model output for the reading times of items with overt pronouns in the fifth window for the experimental and control groups without *order*. 166
D.2 Model output for the reading times of items with overt pronouns in the fifth window for the experimental with *placement* as a simple effect. 166
D.3 Model outputs for the reading times of items with overt pronouns in the fifth window in each half of the experiment. 166

List of abbreviations

1 First Person
2 Second Person
3 Third Person
ACC Accusative
AIC Akaike Information Criterion
ATH Activation Threshold Hypothesis
Att Attachment
BLP Bilingual Language Profile
CEFR Common European Framework of Reference for Languages
CL Clitic
CP Complementiser Phrase
D Determiner
DO Direct Object
DOM Differential Object Marking
DP Determiner Phrase
Emb Embedded
F Feminine
FH Feature Hierarchy
FP First Pass
FUT Future
GP Go Past
HA High Attachment
HS Heritage Speaker
IH Interface Hypothesis
IP Inflectional Phrase
IQR Interquartile Range
L1 First Language
L2 Second Language
L2er Second Language Speaker
LA Low Attachment
LC Late Closure
LF Logical Form
LoI Length Of Instruction
LoR Length Of Residence
M Masculine
Mat Matrix
DAT Dative

GEN Genitive

NOM Nominative

NP Noun Phrase

OQPT Oxford Quick Placement Test

PAH Position of Antecedent Hypothesis, also referred to as the Position of Antecedent Strategy (PAS)

PF Phonological Form

PL Plural

PR Pseudorelative

Pred Predicate

POS Possessive

PRS Present

PRG Progressive

PRFH PR-First Hypothesis

Pro Pronoun

PST Past

RC Relative Clause

REFL Reflexive

SG Singular

SpecIP Specifier of the Inflectional Phrase

SVO Subject-verb-object

T Tense

TS Topic Shift

TT Total Time

VOT Voice Onset Time

VP Verb Phrase

1 Introduction

1.1 Introduction

In everyday situations, we often encounter utterances that are at least temporarily ambiguous between multiple readings. As a concrete example, take the Italian sentence in (1). That sentence is globally ambiguous as the pronoun in the embedded CP (*lei* or *ø*) may be interpreted as co-referent with the matrix subject (i.e., the elderly woman), the matrix object (i.e., the girl), or some unmentioned third referent (i.e., someone else). Barring further information (linguistic or contextual), the hearer/reader would not be able to exclude any of these possible readings. Despite this global ambiguity, native speakers do not exhibit indeterminacy in their interpretations of such sentences. Rather, they have repeatedly been found to exhibit different biases in both their offline interpretations and their online processing depending on the pronominal form (i.e., *lei* or *ø*).

(1) a. *La anziana signora saluta la ragazza mentre (lei / ø) attraversa la strada.* [Italian]
the elderly woman greets the girl while she *pro* crosses the street

'The elderly woman greets the girl while she crosses the street.'

This has led to considerable interest in how speakers not only acquire these forms and various related properties (e.g., the general availability of postverbal subjects, Rizzi 1982), but also how their interpretive biases are acquired. As a result, pronominal subjects have been extensively studied in various populations and from different theoretical backgrounds with a particular wealth of research focusing on speakers of 'null-subject' languages like Italian that allow for phonologically overt (e.g., *lei*) and null (i.e., *ø*) subject pronouns. In general, previous work has found that even when the syntactic optionality between null and overt subjects in those languages has been acquired, acquirers' production and interpretations of such forms may exhibit protracted differences from monolingual adults, particularly in the case of bilingual individuals.

Focusing on Italian as a concrete example of the relevant languages, monolingual children have been found to exhibit behaviour that is consistent with them having acquired the underlying syntactic optionality of null and overt subjects from the earliest stages of production. In particular, (typically developing) child acquirers of Italian do not go through a 'root infinitive' stage in which they produce uninflected main verbs as is common for acquirers of non-null-subject languages such as English or French (e.g., Hyams 1986, Pierce 1992, Rizzi 1993, Wexler 1994). Rather, child acquirers of Italian appear to have mastered verbal agreement even by their earliest production (Guasti 1993). Additionally, Italian-speaking children have been found to omit subjects noticeably more often than their English (i.e., non-null-subject) speaking peers (Valian 1991), with convergent results also reported for the acquirers of other null-subject languages such as (Brazilian) Portuguese (Valian & Eisenberg 1996). Within child acquirers of Italian, the realisation of null and overt subjects has also been observed to be sensitive to argument structure. Although children generally produce utterances with null subjects in line with adults, both groups are noticeably more likely to produce an overt subject when it surfaces postverbally with an unaccusative predicate (Lorusso, Caprin & Guasti 2004). Given that null subjects are argued to occupy a particular (preverbal) position (Cardinaletti 2004), and given that the subject of unaccusative predicates may remain *in situ*, that is, in its base generated postverbal position (Belletti 1988), this has been taken as evidence that even very young children (1;6 – 3;0) are aware of the structural restrictions on null subjects. Despite the non-trivial evidence that the underlying syntactic optionality between null and overt subjects along with its structural restrictions is acquired very early, children's interpretation (although not necessarily production, e.g., Lorusso et al. 2004) of such forms exhibits protracted non-convergence with adult norms; child acquirers of Italian have been found to differ from adults in their interpretation of overt pronouns even after entering the school system (6;2 – 7;11, Sorace, Serratrice, Filiaci & Baldo 2009; 6:0 – 8;9, Vogelzang, Guasti, Rijn & Hendriks 2021) with related results also reported for (Mexican) Spanish,

another null-subject language (Shin & Cairns 2012). Convergent results have also been reported for postverbal subjects in narrow new-information focus (i.e., in response to a 'who'-question), the availability of which is related to the availability of null pronouns (Rizzi 1982, Belletti 2001, 2004). Although school-aged Italian children have clearly acquired that such postverbal subjects exist in their target variety, their interpretation of such forms has not yet fully converged with that of adult native speakers (6;1 – 7;4, Cairncross & Dal Pozzo 2021).

In a similar vein, work with bilingual children (i.e., those with two first languages, L1s, or early acquirers of a second language, L2) acquiring a null-subject language has repeatedly shown that although those children are aware of the availability of null subjects in their relevant language(s) (i.e., they produce null pronouns where appropriate), they overproduce overt pronouns in their null-subject language relative to their monolingual peers (e.g., Paradis & Navarro 2003, Serratrice, Sorace & Paoli 2004, Pinto 2006, Hacohen & Schaeffer 2007, Silva-Corvalán 2014, Torregrossa, Bongartz & Tsimpli 2019). Similarly, they have also been found to exhibit higher rates of non-target-like interpretations for the same forms (e.g., Serratrice 2007, Sorace et al. 2009). While many of the previous studies have focused on children acquiring a null-subject language alongside a non-null-subject one such as English, the reported effects cannot be entirely attributed to the presence of the non-null-subject language. This is because children's performance is affected by non-linguistic factors such as the dominant language of their community (e.g., Italian-English bilinguals in the UK vs in Italy Sorace et al. 2009). Moreover, higher rates of non-target-like interpretations for overt pronouns have also been observed in bilingual children who speak two null-subject languages (Sorace et al. 2009, although see the discussion of Filiaci, Sorace & Carreiras 2014 i.a. in section 2.2).

Subject pronouns in null-subject languages have also been studied extensively in 'heritage speakers (HSs).' To some extent, HSs overlap with early bilinguals in that they acquire both the dominant language of their community as well as a minority language early in life. However, research on HSs is also interested in later stages of development and generally focuses on young adults. Taking Spanish in the United States as a concrete example, heritage speakers of Spanish have been found to exhibit related differences from monolingual speakers of their null-subject language in the production (e.g., Montrul 2004, Otheguy, Zentella & Livert 2007), interpretation (e.g., Keating, VanPatten & Jegerski 2011) and online processing (e.g., Keating, Jegerski & Vanpatten 2016) of pronominal forms, suggesting that early bilinguals' interpretation and production of such forms may never fully converge with the monolingual target.

Turning to adult L2 acquisition, while much of the early work largely debated whether native speakers of a non-null-subject language who are acquiring a null-subject L2 may fully acquire the availability of null subjects (i.e., parameter resetting in the principles and parameters framework) and vice versa (e.g., White 1985, Hilles 1986, Liceras 1989, Tsimpli & Roussou 1991, Liceras & Díaz 1999), subsequent work has focused on how such forms are produced (e.g., Pérez-Leroux & Glass 1999, Montrul & Louro 2006, Belletti, Bennati & Sorace 2007, Rothman 2009), interpreted (e.g., Pérez-Leroux & Glass 1999, Gurel 2006, Sorace & Filiaci 2006, Belletti et al. 2007, Roberts, Gullberg & Indefrey 2008, Rothman 2009, Jegerski, VanPatten & Keating 2011, Di Domenico, Baroncini & Capotorti 2020), and processed (e.g., Roberts et al. 2008, Bel, Sagarra, Comínguez & García-Alcaraz 2016, Cunnings, Fotiadou & Tsimpli 2016) by L2 speakers. Focusing on studies in which the L2 is the null-subject language, many studies have found that even though such speakers have acquired and are able to correctly produce and interpret null subjects, they additionally exhibit protracted non-convergence with native speakers of their target language with regard to the production and interpretation of overt pronouns. However again, these remnant difficulties likely should not be attributed entirely to the presence of the non-null-subject language as the overproduction of overt pronominal subjects has also been attested in adult speakers whose L1s are typologically similar null-subject languages (e.g., Lozano 2006, Margaza & Bel 2006, Contemori & Dussias 2020) with complementary non-convergence also reported in their interpretation of such forms (e.g., Lozano 2017). Likewise, if we consider the position of subjects relative to information status, L1-English and German L2 speakers of Italian, similar to monolingual children, have been found to have acquired the

availability of postverbal subjects in the relevant context while also exhibiting protracted non-convergence with native Italian speakers in their production of such forms (Belletti & Leonini 2004, Belletti et al. 2007). Thus synthesising the above, even when there is evidence that the underlying syntactic options relating to the availability of null subjects have been acquired, learners – and in particular bilingual learners – may still experience residual difficulties in their production and interpretation of the relevant forms.

More germane to the present thesis, there is also a growing recognition that even once these forms and their relevant interpretive biases are fully acquired, they are not immutable. Rather, aspects of an individual speaker's L1 may be affected by the acquisition and use of an L2 in adulthood, a phenomenon generally referred to as *L1 attrition*. To date, attrition has been studied in a variety of linguistic domains with the clearest examples being its effects on phonemic realisation and lexical access/diversity. As an example of the former, native speakers of French have been observed to exhibit a shorter voice onset time (VOT) for vowels following plosives (e.g., /t/) than native English speakers. However, after the acquisition and prolonged immersion in their L2 English (mean = 12.2 years), native French speakers have been observed to exhibit a longer VOT in their L1 French than monolingual speakers. The opposite pattern has also been observed in L1-English speakers who had acquired and experienced prolonged immersion in their L2 French (mean = 11.7 years), suggesting that the acquisition of an L2 may lead to changes to the realisation of phonemes in the L1 (Flege 1987). Since that original study, various authors have argued that attrition may lead to assimilation effects (e.g., Flege & Hillenbrand 1984, Major 1992, Mayr, Price & Mennen 2012) as well as dissimilation effects (e.g., Flege & Eefting 1987) at both the segmental and suprasegmental levels. Moreover, while many previous studies have focused on L1 speakers with relatively long periods of immersion in their L2, there is some evidence from L1-English L2-Korean speakers that changes to phonemic realisation may surface very early, after only a few weeks of exposure to the L2 (Chang 2012). Turning to the lexicon, it has been often observed that adults who grew up speaking more than one language may have smaller vocabularies in their L1 and experience difficulties retrieving lexical items relative to monolingual controls (e.g., picture naming: Gollan, Montoya, Fennema-Notestine & Morris 2005, Bialystok, Craik & Luk 2008, Gollan, Montoya, Cera & Sandoval 2008, Ivanova & Costa 2008; verbal fluency: Bialystok et al. 2008; perceptive and expressive vocabulary Portocarrero, Burright & Donovick 2007; tip of the tongue states: Ecker 2004). Related effects may also be observed in the L1 of adult L2 acquirers (e.g., lexical diversity in free speech: Schmid & Dusseldorp 2010, Yilmaz & Schmid 2012, Schmid & Jarvis 2014; verbal fluency: Schmid & Dusseldorp 2010; picture naming: Schmid & Yilmaz 2021). Moreover, there is again some evidence that changes to lexical access may surface after only a few months of immersion in one's L2 (Linck, Kroll & Sunderman 2009, Baus, Costa & Carreiras 2013). These results highlight that our L1s remain more plastic than is sometimes assumed and even limited contact with an L2 may result in measurable divergence from monolingual norms. For recent overviews of how attrition may affect other domains, see Schmid & Köpke (2019), although note that some of the authors therein assume definitions of attrition which differ slightly from that which we will adopt in section 2.3.

Within the present thesis, we will primarily focus on how attrition affects the interpretation of more syntactic phenomena. Specifically, we start by focusing on the interpretation of null and overt pronouns which has been observed to be affected by attrition in a variety of null-subject languages (e.g., Bulgarian: Köpke & Genevska-Hanke 2018; Italian: Tsimpli et al. 2004; Spanish: Chamorro, Sorace & Sturt 2015a, Martín-Villena 2023; Turkish: Gürel 2004, Gürel & Yılmaz 2011) and is often reported to result in an overextension of overt pronouns similar to the pattern reported for the other bilingual populations mentioned above. The present thesis will then also incorporate 'RC' attachment ambiguities into the discussion. An example of this kind of ambiguity is given in (2). Such sentences are globally ambiguous because the embedded 'RC' (indicated by brackets) may attach to either of the preceding non-subject DPs (i.e., the 'doctor' or the 'son of the doctor'). As with the interpretation of null and overt pronouns, however, (monolingual) native speakers do not exhibit indeterminacy in their interpretation of such strings, instead relying on structural parsing biases. Despite

limited interest in attrition and attachment ambiguities, there are some previous results from Spanish that suggest these ambiguities may also be affected (Dussias 2003, 2004, Dussias & Sagarra 2007). However, it is unclear how we can interpret and integrate those findings into a more general account of attrition in light of recent theoretical developments (i.e. the PR-First Hypothesis, Grillo 2012, Grillo & Costa 2014).

(2) *Gianni ha visto il figlio del medico [che correva].* [Italian]
Gianni has seen the son of.the doctor that ran

'Gianni saw the son of the doctor (that was) running.'

Nonetheless, we believe the comparison of these phenomena under attrition to be useful for two reasons. First, by comparing different types of interpretive biases, we can tease apart what makes certain biases vulnerable to attrite. Second, a better understanding of how interpretive biases are affected by attrition is useful for the study of bilingualism more generally. This is because the study of attrition allows us to isolate the effects of being multilingual from the acquisition of the phenomenon under investigation. That is to say, given we trivially know that an attriter (i.e., those undergoing attrition) fully acquired the relevant phenomena before becoming bilingual, any deviance from monolingual norms cannot be accounted for in acquisitional terms. Rather, they should be accounted for in terms of how a speaker's languages interact in the mind. As such, the study of attrition can better inform us about the reasons for non-convergence in other bilingual populations (e.g., L2 pronominal resolution in null-subject languages).

1.2 Originality

We take the present thesis to be an original contribution to the literature for several reasons. First, in our review of the previous work on syntactic attrition, we will argue that other authors' results exhibit two relevant patterns. On the one hand, there is evidence from some studies for the often-discussed relaxation of discourse constraints for specific forms (e.g., overt pronouns in null-subject languages). On the other hand, we will argue that previous results also indicate a more general tendency, which we believe has so far gone largely unnoticed. That is to say, we will argue that the now bilingual L1 speakers (i.e., potential attriters) in the same earlier studies also exhibited an increased reliance on interpretive biases already present in their L1s when faced with ambiguity.

Second, by re-using pronominal items from the previous literature on attrition, we will demonstrate that a notable example of the increased reliance on interpretive biases (i.e., the increased subject bias for null pronouns reported by Tsimpli et al. 2004) is replicable, thereby increasing the credence that this effect should be taken seriously and accounted for.

Third, the present thesis also adds to the limited body of work exploring the attrition of intrasentential pronominal resolution from an online perspective. Results indicate that even the prime example of a syntax-discourse interface phenomenon (i.e., overt pronouns in an Italian-style null subject language) which has been argued to undergo a relaxation of its discursive constraints under attrition may also instead be affected by the general tendency to increase one's interpretive biases resulting in more – not less – monolingual-like processing *contra* Chamorro et al. (2015a). We will argue that this is not without precedent and that previous results for overt pronouns paint a conflicting picture.

Forth, in order to further explore the novel pattern mentioned above, the present thesis will also study a phenomenon which has so far received limited attention from researchers investigating attrition, so-called 'RC' attachment ambiguities. In doing so, our results from the interpretation of globally ambiguous items will extend previous findings from L1-Spanish, L2-English attriters (i.e., those undergoing attrition) (Dussias 2003) to a new language pair (i.e., L1-Italian, L2-English). We take this finding to indicate that the observed pattern is a general one which should be accounted for.

Fifth, we believe the present thesis to be the first study to investigate the effects of attrition on 'RC' attachment biases in light of Grillo & Costa's (2014) PR-First Hypothesis. Once the previously overlooked distinction between RCs and PRs is taken into consideration (although string identical, the two are structurally distinct), the offline interpretation study in the present thesis leads us to the opposite conclusion from that in Dussias (2003, 2004) and Dussias & Sagarra (2007). We will argue that attrition of parser biases is not due to a loss of a bias in the L1 under pressure from a conflicting bias in the L2. Rather we will argue that attrition results in an increased reliance on another bias already present in the L1. As such, we take that study to provide an additional example of the general tendency under attrition to rely more heavily on interpretive biases argued for in the literature review.

Sixth, given the limited previous work on the online processing of PR/RC parsing ambiguities even with monolingual speakers, the present thesis also evaluates the online predictions of the PR account. This is done to establish a baseline against which to compare potential attrition effects. However, in doing so, results from the control participants provide two new pieces of evidence in support of the PR account. On the one hand, we extend previous results from Spanish (Aguilar, Ferré, Gavilán, Hinojosa & Demestre 2021) to Italian, showing that native speakers are sensitive to the PR/RC distinction in online parsing and that this, in turn, affects attachment biases. On the other hand, we provide a novel argument in support of the PR account by demonstrating replicable interference effects in final interpretations of temporarily ambiguous items. We will argue that these interference effects are consistent with the idea that PR parses are preferred over RC parses *ceteris paribus* (PR-First Hypothesis).

Seventh, by studying the potential effects of attrition on both syntactic parsing ambiguities and intrasentential pronominal resolution within the same participant using the same tasks, the present thesis presents a novel comparison allowing us to partially disentangle the role of biases (interpretive or processing) from the syntax-discourse interface (Sorace & Filiaci 2006, Sorace 2011). This is because following previous research we will assume the interpretation and online processing of pronominal subjects in Italian to sit at the interface between syntax and discourse whereas we will assume (pseudo-)relative clause attachment ambiguities are conditioned by principles of computational efficiency at the level of the parser. Therefore, should there be something about more external interface structures which causes them to be more vulnerable to attrite as has been previously suggested, we would expect attrition to affect the two phenomena differently. However, should attrition lead to an increased reliance on certain biases regardless of interface status as will be argued for in our review of the previous literature, we might expect the two phenomena to pattern together.

1.3 Thesis structure

The present thesis is organised in the following way. Chapter 2 starts by describing the first core phenomenon for the present thesis, namely pronominal resolution in Italian and related null-subject languages. After establishing a working definition of attrition for the present thesis, that chapter introduces the Interface Hypothesis (IH) which has been an influential framework within the study of attrition. The chapter then presents a review of previous literature which mostly focuses on aspects of null-subject languages and raises a very general starting question for the thesis regarding interpretive biases for other types of structures. To further explore that question, the chapter introduces the second core phenomenon explored in the present thesis, 'RC' attachment ambiguities, and lays out a series of testable sub-questions to guide the rest of the thesis.

Chapter 3 focuses on how attrition affects offline interpretive biases in the two phenomena of interest. To that end, the chapter presents a sentence interpretation task in which participants – both Italians in their home community and Italians living abroad – were asked to identify their preferred interpretation of globally ambiguous stimuli.

So that we could also investigate how attrition affects biases in online processing as has become the focus of recent work (e.g., Sorace 2011), chapter 4 motivates the need for the construction of two sets of novel, locally ambiguous stimuli (i.e., one for pronominal ambiguities and one for attachment ambiguities). That chapter also details a two-step evaluation process for our novel stimuli in which we asked native linguists to judge our items and then conducted a norming study with non-linguists to identify any items potentially biased by factors (e.g., plausibility) other than our specific experimental manipulations.

Chapters 5 and 6 each present three experiments which explore biases in the online processing of pronominal and attachment ambiguities respectively. In both chapters, the participants for the first two experiments (a self-paced reading study, and an eye-tracking-while-reading study) consisted exclusively of native Italian speakers still residing in Italy. This was done to establish and validate a baseline against which to compare potential attriters. This was particularly relevant for the attachment ambiguities explored in chapter 6 given the limited previous work on the processing of such structures as discussed in chapter 2. In the third experiment, the results from the first (self-paced reading) experiments are compared against a group of native Italian speakers living in an second language (L2) environment whose language background suggests they are a suitable group in which to explore attrition effects.

Tying the experimental chapters together, chapter 7 presents a general discussion in which we (i) compare the online results from chapters 5 and 6, (ii) discuss the similarities and differences between the effects of attrition on online and offline biases, and (iii) integrate our findings into the larger literature.

2 General background and research questions

2.1 Introduction

The present chapter is divided into four main parts. Section 2.2 begins with a description of the relevant aspects of the first core phenomenon explored in this thesis, namely the interpretative biases for pronominal subjects in null-subject languages like Italian. Section 2.3 then introduces L1 attrition with a particular focus on the IH as well as structures sitting at the so-called 'syntax-discourse interface.' A review of the previous literature generally taken in support of the IH will then raise a very general starting question for the rest of the thesis regarding interpretive biases beyond the usually studied properties of null-subject languages. Building off of this, section 2.4 will introduce the second core linguistic phenomenon investigated in the thesis, namely 'RC' attachment ambiguities. Integrating attachment ambiguities into the preceding section will lead to the formulation of the specific research questions tackled in the following chapters (section 2.5).

2.2 Pronominal subjects in Italian (and related languages)

Italian is a null-subject language in which null subjects are licensed by agreement on T (Rizzi 1982, 1986). As such, the subject of a well-formed finite clause may be phonologically overt or null (3a). This contrasts with non-null subject languages like English, where null subjects are restricted (3b).

(3) a. *(Rita e Maciej / ø) amano i gatti.* [Italian]
Rita and Maciej *pro* love the cats

'(Rita and Maciej / they) love cats.'

b. (Rita and Maciej / *ø) love cats.

Despite this syntactic optionality, null and overt pronominal forms in Italian are not in true free variation. Rather, they are constrained by discourse factors such as topicality and focus with overt pronouns representing the 'marked' option (Cardinaletti & Starke 1999). When the referent of the subject represents a shift in the topic, native Italian speakers preferentially produce an overt pronoun. Conversely, when the referent of the subject has already been introduced and is topical, the same speakers preferentially produce a null pronoun. This asymmetry extends to comprehension. As such, in an intrasentential context like that in (4) in which the matrix subject is generally assumed to be the default topic (Reinhart 1981), null pronouns are preferentially interpreted as coreferential with the matrix subject, whereas overt pronouns are preferentially interpreted as picking out some other referent (i.e., the object referent or a third unspecified referent).

(4) *Gianni$_i$ saluta Marco$_j$ quando (ø$_{i/?j/?k}$) / lui$_{?i/j/k}$) attraversa la strada.* [Italian]
Gianni greets Marco when *pro* he crosses the road

'Gianni greets Marco when he crosses the road.'

To account for the observed interpretive biases for null and overt pronouns in intrasentential contexts, Carminati (2002) proposed the Position of Antecedent Hypothesis (PAH). The PAH posits that in online processing speakers of Italian exhibit two complementary biases. On the one hand, speakers exhibit a bias

for null pronouns to corefer with the potential antecedent that occupies the highest SpecIP[1] position. This is generally – but not always – the subject as in (5). On the other hand, it posits that there is a bias for overt pronouns to be coreferential with a possible antecedent in a structurally lower position, for example, the direct object in (5).

(5) *Quando Mario ha telefonato a Giovanni, (lui / ø) aveva appena finito di mangiare.* [Italian]
when Mario has called to Giovani he *pro* had just finished of eat

'When Mario called Giovanni, he had just finished eating.'

To evaluate the predictions of the PAH, Carminati first presented native Italian speakers living in Italy (N = 40) with a (non-cumulative) self-paced reading task. For this task, they constructed a series of biclausal sentences in which the embedded CP preceded the matrix CP that contained the ambiguous pronoun (null vs overt). Pronominal coreference was manipulated pragmatically in the second CP (SpecIP vs non-SpecIP). After each item, participants were additionally presented with a comprehension question. An example item is presented in (6) with vertical bars to indicate the window boundaries along with its corresponding comprehension question in (7). Consistent with the PAH, Carminati found that the matrix CP was read more slowly when null subjects were pragmatically biased to be coreferential with the embedded object than when they were biased toward the embedded subject. For overt pronouns, they observed the opposite pattern. Under statistical analysis (ANOVA), this resulted in a significant interaction of *pronoun* by *antecedent*.[2] A convergent trend and significant interaction were observed in response times to the comprehension question. Moreover, the PAH was reflected in response accuracy in their pragmatically disambiguated items (null: SpecIP: 88.7%, nonSpecIP: 70.4%; overt: SpecIP: 80.4%, non-SpecIP: 89.1%). In that case, Carminati (2002) not only reports a significant interaction but also a significant *post hoc* *t*-test indicating that PAH violations resulted in lower accuracy for null pronouns.

(6) a. *Dopo che Giovanni ha meso in imbarazzo Giorgio di fronte a tutti, | (ø / lui) si*
After that Giovanni has put in embarrassment Giorgio of front to all, *pro* he REFL
è scusato ripetutamente. [Italian]
is excused repeatedly.

'After Giovanni embarrassed Giorgio in front of everyone, he apologised repeatedly.'

b. *Dopo che Giovanni ha meso in imbarazzo Giorgio di fronte a tutti, | (ø / lui) si*
After that Giovanni has put in embarrassment Giorgio of front to all, *pro* he REFL
è offeso tremendamente.
is offended tremendously

'After Giovanni embarrassed Giorgio in front of everyone, he was tremendously offended.'

(7) *Chi si è scusato?* [Italian]
who REFL is excused

'Who apologised?'

a. Giorgio

b. Giovanni

[1] Within modern generativist work, clauses are generally divided into three domains: the verb phrase (VP), the inflectional phrase (IP), and the complementiser phrase (CP). The VP encodes thematic information (i.e., who did what to whom), whereas the IP anchors the event encoded by the VP and the CP generally encodes discourse-related information.

[2] *Post hoc*s to further explore whether the effect of *antecedent* was significant within each pronominal type were not reported.

Building off of this, Carminati also demonstrated quantitatively that the PAH affects offline interpretations in fully ambiguous contexts. To show this, they presented native Italian speakers (N = 44) with globally ambiguous items as in (8). As in their earlier experiment, the embedded CP always preceded the matrix CP which contained the ambiguous pronoun. Under each item, participants were presented with the two possible interpretations of the embedded CP from the linguistic context and asked to select their preferred interpretation. Responses indicated clear interpretive biases in line with the predictions of the PAH. That is to say, when presented with a null pronoun, participants overwhelmingly selected the SpecIP-antecedent reading (80.72%). Conversely, when presented with an overt pronoun, they overwhelmingly selected the non-SpecIP-antecedent reading (83.33%). This resulted in a significant interaction of *pronoun* by *antecedent*.

(8) *Marta scriveva frequentamente a Piera quando (ø / lei) era negli Stati Uniti.* [Italian]
Marta wrote frequently to Piera when *pro* she was in.the states united

'Marta wrote to Piera frequently when she was in the United States.'

a. *Quando Marta era negli Stati Uniti.*

b. *Quando Piera era negli Stati Uniti.*
when M/P was in.the states united

'When M/P was in the United States.'

Furthermore, in a series of similar self-paced reading experiments with native Italian speakers (Ns between 42 and 44), Carminati (2005) showed that even when such sentences are fully disambiguated via grammatical features (9) as opposed to plausibility, a significant PAH effect is observed with null pronouns in Italian (In those experiments, they did not consider overt pronouns.).

(9) *Quando Maria cerca Roberto, | ø diventa ansios(-a/-o)* [Italian]
when Maria.F looks.for Roberto.M *pro* becomes nervous-F/-M

'When Maria looks for Roberto, (s/he) becomes nervous.'

However, Carminati also found that the type of disambiguation is relevant. Specifically, they were interested in the position of features within the Feature Hierarchy (FH) presented in (10).[3] This hierarchy did not derive from the psycholinguistic literature but instead was based on the frequency of (co-)occurrence noted in topological literature (Greenberg 1963) and the fact that some syntactic phenomena such as split ergativity appear to be sensitive to the pattern (Silverstein 1985). The FH implies that if a language has a given feature on the hierarchy, it must also have the features to the left (e.g., if a language has gender, it must also have number and person).

(10) Feature Hierarchy: Person > Number > Gender

Nonetheless, following Harley & Ritter's (2002) idea that the hierarchical organisation of these features relates to their 'cognitive significance,' Carminati (2005) hypothesised that there is a correlation between the position of a feature in the FH and its disambiguating power, namely the higher the feature sits, the better it should be at disambiguating stimuli. With respect to intrasentential pronominal resolution, they predicted that the penalty for violating the PAH should be smaller for disambiguating features that sit higher in the FH. Consistent with this prediction, Carminati observed that reading times for items in which null pronouns were forced toward a non-SpecIP antecedent were slower when disambiguated via gender (11a) than when they

[3]Subhierarchies within each feature type have also been proposed (e.g., Silverstein 1985, Harley & Ritter 2002).

were disambiguated by number (11b). That is to say, the cost of violating the SpecIP bias for null pronouns was smaller in the number disambiguated items. Although the original author does not appear to have run an ANOVA to explore the interaction between *feature* and *antecedent*, planned pairwise comparisons indicated that indeed PAH violations were significantly more costly for gender disambiguated items. Moreover, a similar effect was observed with items disambiguated by person and number (12a), which were read even faster than items disambiguated by number alone (12b).

(11) a. *Quando Maria lo cerca, | ø diventa ansios-o.* [Italian]
when Maria.F CL.M looks.for *pro* becomes anxious-M

'When Maria looks for him, he becomes anxious.'

b. *Quando i Rossi lo cercano, | ø diventa ansios-o.*
when the Rossi.PL CL.SG looks.for *pro* becomes anxious-SG

'When the Rossis looks for him, he becomes anxious.'

(12) a. *Quando (ho/hai) salutato i nonni, | ø sembravano veramente tristi.* [Italian]
when have(.1/.2).SG greeted the grandpar.3.PL *pro* seemed.3.PL truly sad.

'When (I/you) said good by to the grandparents, they seemed truly sad.'

b. *Quando Gianni ha salutato i nonni, | ø sembravano veramente tristi.*
when Gianni.3.SG has greeted the grandpar.3.PL *pro* seemed.3.PL truly sad.

'When Gianni said good by to the grandparents, they seemed truly sad.'

Since Carminati (2002), the PAH has also been tested in a variety of null-subject languages with results indicating some similarities and dissimilarities between even closely related languages (e.g., Italian and Spanish). Working on (Iberian) Spanish, Alonso-Ovalle, Fernández-Solera, Frazier & Clifton (2002) presented participants (N = 80) with ambiguous items as in (13).

(13) *Juan pegó a Pedro. (ø / El) está enfadado.* [Spanish]
Juan hit DOM Pedro pro he is angry

'Juan hit Pedro. He is angry.'

After each item, the participants were asked a comprehension question eliciting the subject of the second CP (e.g., *¿Quién está enfadado?* - 'Who is angry?'). Results indicated that participants preferentially interpreted null subjects as being coreferential with the subject of the preceding CP (73.2%). Conversely, overt pronouns were interpreted at chance (50.2% subject). On the one hand, this result suggests that (Iberian) Spanish is consistent with the PAH at least with regard to null pronouns. On the other hand, these results suggest that there may be relevant microvariation between even closely related null-subject languages. Recall that working with (intrasentential) ambiguous pronominal forms, Carminati (2002) found a clear non-subject bias for overt pronouns in Italian (83.33%). This clearly contrasts with the absence of an antecedent bias reported for the (Iberian) Spanish speakers, a point to be kept in mind when considering the interpretation of Spanish-Italian bilingual data.

Extending this to processing and intrasentential contexts, Filiaci et al. (2014) translated Carminati's (2002) plausibility disambiguated items (6) into Spanish. An example sentence in both languages is provided in (14) with vertical bars to indicate window boundaries.

(14) a. *Dopo che Giovanni ha criticato Bruno così ingiustamente,* | *(ø / lui) si è scusato*
after that Giovanni has criticised Bruno so unjustly | pro he REFL is excused
ripetutamente. [Italian]
repeatedly

'After Giovanni embarrassed Bruno so unjustly, he apologised repeatedly.'

b. *Después de que Bernardo critcó a Carlos tan injustamente,* | *(ø / el) le pidió*
after of that Bernardo criticised DOM Carlos so unjustly | pro he CL asked.for
disculpas. [Spanish]
forgiveness

'After Bernardo criticised Carlos so unjustly, he apologised.'

The authors then presented otherwise identical self-paced reading experiments to a group of native Italian speakers (N = 32) and a group of native (Iberian) Spanish speakers (N = 32). Reading times for the second CP indicated that both groups exhibited a similar SpecIP bias for null pronouns. For overt pronouns, however, although the Italian group exhibited the expected PAH effect with items with SpecIP antecedents being read more slowly than items with object antecedents, the Spanish group did not differ in their reading times for the two antecedents. Thus, similar to the offline results in Alonso-Ovalle et al. (2002), the (Iberian) Spanish data for null pronouns is consistent with what has been previously reported for Italian, while results from items with overt pronouns suggest that speakers of (Iberian) Spanish are more flexible with / lack a non-SpecIP antecedent bias. For a recent review of the literature comparing Italian and Spanish with respect to the PAH, please see Contemori & Di Domenico (2021).

As a final point on the PAH that will become relevant, there is evidence that the PAH is operative – but weaker – in contexts containing a single textually given possible antecedent. To explore the possibility, Carminati (2002) presented native Italian speakers (N = 44) with ambiguous items with a single textually-given possible referent as in (15).

(15) *Gregorio ha detto che (ø / lui) sarà presente al matrimonio di Maria.* [Italian]
Gregorio has said that pro he BE.FUT present at.the wedding of Maria

'Gregorio said that he will be at Maria's wedding.'

a. *Gregorio stesso sará presente al matrimonio.*
Gregorio same BE.FUT present at.the wedding

'Gregorio himself will be at the wedding.'

b. *Una persona diversa da Gregorio sarà presente al matrimonio.*
a person different from Gregorio BE.FUT present at.the wedding

'A person different from Gregorio will be at the wedding.'

In such contexts, native speakers overwhelmingly preferred SpecIP antecedent readings over the new referent option, however, the bias was significantly stronger for null pronouns (96.59%) than overt ones (85.79%). This result is notable for two reasons. First, it underlines that even in these highly reduced linguistic contexts (at least in matrix-first items), new-referent, or obviative, readings are a last resort option. Nonetheless, despite the strong tendency to avoid obviative readings, a residue of the PAH remains. Namely, the SpecIP-antecedent bias for null pronouns is more pronounced than for overt pronouns in such contexts. The original author took this to suggest that while the first half of the PAH (that null pronouns prefer an

antecedent in SpecIP) is unaffected by context, the second part of it (that overt pronouns prefer an antecedent in a lower position) is more flexible and sensitive to contextual factors.

Given the success of the PAH, it has been widely adopted by work on the processing and interpretation of pronominal subjects in null-subject languages, as well as work on attrition more specifically (e.g., Chamorro et al. 2015a, Martín-Villena 2023). As such, we will also adopt the PAH in the present thesis. However, before discussing how these biases may be affected by attrition, we should first provide a working definition for attrition itself.

2.3 Attrition

2.3.1 A working definition

As the term 'attrition' has been used to mean different things by different authors, it is worth setting out how the term will be used in the present thesis. Here, we will adopt the working definition that attrition refers to any changes within an individual speaker's L1 due to the acquisition and use of a L2 after the L1 had been acquired. As such, attrition is distinct from language change due to pathology as well as language loss/change at the community level in language contact situations (sometimes called 'generational attrition'). Additionally, under our working definition, attrition is distinct from cases of 'interrupted' or 'incomplete' acquisition in which the L1 may not have been fully acquired prior to exposure to the L2 (e.g. in cases of adoption by speakers from another speech community or in the case of heritage speakers, HSs, who acquire their L1 in a minority context). That is not to say that such speakers may not undergo attrition (i.e., we could relativise attrition to the acquisition of specific structures, rather than the L1 as a whole). However, considering potential attrition within HSs introduces additional levels of complexity. This is because such speakers are likely to receive non-trivial L2 input while at least some aspects of their L1s are still developing. Moreover, their L1 target is likely to be qualitatively different from child acquirers in the home community given that their primary linguistic data likely comes from potential attriters and other HSs in a potentially restricted range of registers/contexts (e.g., Rothman 2007, Pires & Rothman 2009). Therefore, for the scope of the present thesis we restrict ourselves to attrition within adult speakers who can be argued to have fully acquired their L1 prior to large-scale contact with an L2.

There are three further aspects of our working definition that are worth highlighting. First, we take attrition to be a property of the individual. As such, we will not consider experimentally induced L2 to L1 interference (e.g., cross-linguistic priming experiments). Rather, we will focus on potential attrition effects in exclusively L1 contexts. Second, we do not make any *a priori* assumptions regarding whether attrition effects hold at the level of representation or processing. We will investigate both, and any relevant effects argued to hold at either level will be considered attrition and potential explanations for them will be explored. Third, we do not presuppose that attrition is the result of direct transfer from the L2. If we are able to detect a reliable change to individual speakers' L1s as a result of the acquisition and use of an L2, that will be taken as an attrition effect regardless of whether the effect can be attributed to a specific property of the L2. That is to say, we will not exclude the possibility that attrition may be the result of cognitive effects related to becoming bilingual. Although we take that this to be a tacit assumption in work by some other researchers (e.g., Sorace 2011, discussed in section 2.3.2), we believe that being explicit in this assumption allows us to identify new potentially interesting patterns, as will become relevant in section 2.3.3.

As it is defined for the present thesis, attrition is typically – although as we will see not necessarily – studied in migration contexts in which the potential attriter (i.e., the person undergoing attrition) has moved out of their L1 community and is immersed in their L2. As a result, they are likely to receive significantly less L1 input (if any at all) as well as significantly more L2 input than if they had stayed in their home community.

2.3.2 The Interface Hypothesis

Within the study of syntactic attrition, the IH has been a particularly influential framework. Originally coined by Sorace & Filiaci (2006), the IH was proposed to account for non-convergent optionality exhibited by advanced late-L2 learners in the production and comprehension of particular structures. Specifically, those authors suggested that syntactic structures that are conditioned by non-syntactic information (i.e., structures that sit at the interface[4] between syntax and other cognitive domains) are more difficult to fully acquire than structures which are not, resulting in protracted optionality even at the highest levels of attainment (i.e., in so-called 'near natives'). Since Sorace & Filiaci (2006), the IH has also been extended to both child bilingual acquisition (e.g., Sorace et al. 2009, Serratrice, Sorace, Filiaci & Baldo 2012) and first language attrition (e.g., Sorace 2011, Chamorro et al. 2015a, Martín-Villena 2023). For the former, the IH is again couched in terms of acquisitional difficulty. For the latter, it is expressed in terms of susceptibility to attrite.

While the original formulation of the IH proposed a dichotomy between the narrow syntax and interface structures, this has since evolved to a gradient distinction between more 'external' and more 'internal' interfaces, where the structures at the former are hypothesized to be harder to acquire/more susceptible to attrite than structures at the latter (Tsimpli & Sorace 2006, although see also White 2011 for discussion on the variability within interfaces in L2 acquisition). So-called 'external' interfaces require the integration of information external to an individual's mental grammar. The prototypical example of a more external interface structure is overt pronouns in null-subject languages which are taken to sit at the syntax-discourse interface given that their felicitous use/interpretation requires the integration of pragmatic or contextual information. As contexts change, hearers/speakers must update their mental representations thereof and integrate that information into their language processing. By contrast, 'internal' interfaces refer to the integration of information that is available within the speaker's grammar that does not require the speaker to consider contextual information. An example of a more internal interface structure might be DOM in a language like Spanish which is conditioned by semantic features such as [± animacy] (on which see Chamorro, Sturt & Sorace 2015b and Montrul, Bhatt & Girju 2015 for somewhat conflicting results).

Under the umbrella of the IH, two general families of accounts have been proposed: representational and processing accounts. To compare the two, let us take overt subject pronouns in null-subject languages like Italian as a concrete example. Within a representational account of attrition, the discursive constraints on overt pronouns would be argued to become underspecified. An early instantiation of this idea can be observed in Tsimpli et al. (2004). Those authors were interested in how L1 speakers of Italian or Greek (another null subject language) who had achieved near-native proficiency in their L2 English interpreted pronominal subjects. Following Chomsky's (1995) distinction between interpretable (i.e., meaningful) and uninterpretable (i.e., only relevant for the morpho-syntactic computation) features, Tsimpli et al. (2004) predicted that attrition only affects interpretable features. To see how this would work, recall that in the relevant null-subject languages the availability of such subject pronouns is licensed by (uninterpretable) agreement features on T. Recall also that overt pronouns in Italian (but also in Greek) are generally produced in TS contexts, whereas null pronouns are produced otherwise with a similar asymmetry holding for comprehension. As such we might assume in line with those authors that overt pronouns bear a [+TS] feature which null pronouns lack (although the opposite has also been proposed, i.e., a [+ topic] feature present only on null pronouns, Hendriks 2014). Moreover, overt pronouns in the speakers' L2 English are assumed to not be specified for [±TS] given the non-null-subject status of that language. As such, the original authors suggested that under attrition overt pronouns in bilingual Italian (and Greek) become optionally underspecified for [±TS] allowing them to appear in non-TS contexts. A feature-based account of this type would neatly predict an overt/null asymmetry. Namely, it would predict only overt pronouns would begin to appear in unexpected contexts (from a monolingual perspective)

[4]Note 'interface' here does not neatly map onto the Phonological Form (PF) and Logical Form (LF) interfaces assumed in minimalism (Chomsky 1995).

due to a relaxation of their feature specifications. Conversely, it would predict that null would not become overextended as they lack the relevant interpretable features to begin with. As we will see below, at least the second half of this asymmetry appears to be correct with mixed results for overt pronouns.

Since then however, various authors have argued against representation accounts of this type based on the observation that L2 speakers of a null-subject language whose L1 is also a null-subject language of a similar type exhibit overextension of overt pronouns (e.g., L1-Spanish, L2-Italian Bini 1993[5] as cited by Sorace 2011; L1-Greek, L2-Spanish Margaza & Bel 2006 and Lozano 2006). Under the assumption that overt pronouns in those languages have a similar featural specification, this is unexpected (although note the discussion of Filiaci 2010, Filiaci et al. 2014 et seq.). Moreover, as we will see in the following section, the overextension of overt pronouns in attrition contexts does not always result in a relaxation of interpretive biases as might be expected if its featural specification had become bleached.

As such, more recent instantiations of the IH have instead focused on potential processing-based explanations. This is based on the idea that the felicitous use of pronominal forms requires rapid, real-time integration of contextual and linguistic information and that bilinguals may be less efficient in this process. Sorace (2011, 2016) suggests that this could potentially relate to bilinguals experiencing additional cognitive strain due to the need to maintain language control (i.e., to avoid interference from the non-target language) or a trade-off between increased inhibitory control and flexibly updating in late second language speakers (L2ers). Regardless of the exact source, under such a processing-couched account, overt pronouns are overextended to compensate for bilinguals' difficulties in integrating multiple sources of information in real time. As for why overt pronouns specifically are affected, Sorace (2011) suggests that the overt form may serve as a 'processing default' (although it is not the linguistic default). For this, they build off of Carminati's (2002) observation that the antecedent bias for null pronouns is fairly consistent whereas the antecedent bias for overt pronouns is generally more flexible in Italian. Recall that Carminati (2002) found that native Italian speakers in Italy rarely admit subject readings for overt pronouns when there are two possible antecedents textually present (16a, 16.67% of trials), whereas they are noticeably more tolerant of PAH violations when there is only a single possible antecedent present in the item (16b, 85.79%). As such, Sorace (2011) suggests that bilinguals' overexertion of overt pronouns is not qualitatively different from monolinguals'. Rather, they suggest that bilingual populations are simply more permissive than monolingual ones in their relaxation of the PAH. Therefore, a processing-couched account of this type would predict that bilinguals will become less sensitive to the discursive restrictions on overt pronouns, despite their underlying representations remaining intact. Moreover, such an account would predict that these effects are more readily detected via online measures (e.g., eye-tracking in temporarily ambiguous items) than offline ones (e.g., untimed grammaticality judgements).

(16) a. *Marta scriveva frequentemente a Piera quando (lei / ø) era negli Stati Uniti.* [Italian]
Marta wrote frequently to Piera when she ø was in.the states united

'Marta wrote to Piera often when she was in the US.'

b. *Gregorio ha detto che (lui / ø) sarà presente al matrimonio di Maria.*
Gregorio has said that he *pro* be.FUT present at.the wedding of Maria

'Gregorio said that he will be at Maria's wedding.' (Carminati 2002)

Given the central role that the subject realisation/interpretation in null-subject languages has played in much of the previous research on attrition and the IH, in the next section we provide a review of these studies. However, in a departure from earlier work on the IH (e.g., Sorace 2011, Chamorro et al. 2015a), we will not focus on the convergence between 'near native' L2ers and L1 attriters. Rather, we will focus exclusively on

[5] Original in Spanish

now-bilingual L1 speakers (i.e., potential attriters). This is because we believe that by focusing on attrition in its own right, we can observe a pattern potentially unique to attrition which appears to have gone unnoticed so far. Namely, we will argue that previous evidence indicates that attrition may lead to increased interpretive biases that are not directly attributable either to transfer (be it at the level of representations or processing) from the L2 or the overextension of overt pronouns as a default form in processing.

2.3.3 Previous work on attrition in null-subject languages

2.3.3.1 Attrition and interpretive biases

To explore their feature-interpretability based predictions, Tsimpli et al. (2004) conducted a picture selection task to tap interpretive biases for intrasentential pronominal resolution in ambiguous items. For this experiment, the authors recruited L1-Greek (N = 19) and L1-Italian (N = 20) speakers who had achieved 'near-native' proficiency in their L2-English and had been immersed in their L2 for a minimum of 6 years. They additionally recruited a control group for each language (Ns = 20). Items consisted of biclausal sentences in which the ambiguous pronominal form always appeared in the embedded CP. In half of their items, the matrix clause preceded the embedded CP (17a). In the other half, the order was reversed (17b).

(17) a. *L'anziana signora saluta la ragazza [quando (ø / lei) attraversa la strada].* [Italian]
the.elderly woman greets the girl when (pro / she) crosses the road

'The elderly woman greets the girl when she crosses the road.'

b. *[Mentre (ø / lei) esce dall'ascensore], l'infermiera urta la donna delle pulizie.*
while (pro / she) exits from.the.elevator the.nurse bumps.into the woman of.the cleaning

'While she exits the elevator, the nurse bumps into the cleaning lady.'

Experimental sentences were presented to participants in their L1s along with three pictures corresponding to the logically possible interpretations of the ambiguous pronoun: coreferential with the matrix subject, coreferential with the matrix object, or some unspecified new-referent. An example of the three images for the item in (17a) is provided in Figure 2.1. Participants were invited to select any and all images that corresponded to the given item. In fact, the authors explicitly encouraged selecting more than one image per item. The reported translation of their instructions states that '[i]n some cases only one picture matches the meaning of the sentence, in other cases, there can be more than one' (Tsimpli et al. 2004: pp. 276-277).

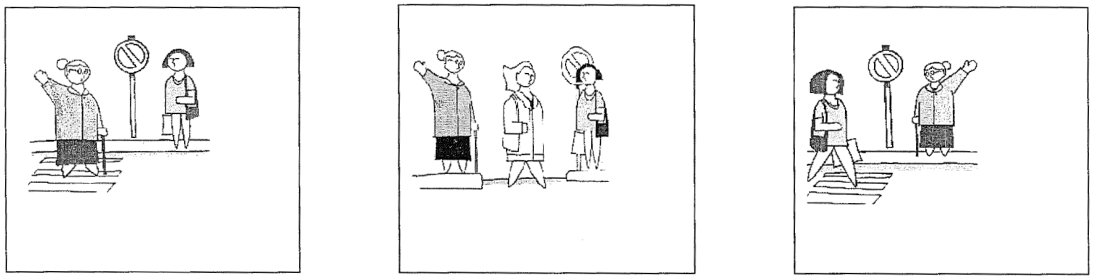

(a) matrix-subject reading (b) new-referent reading (c) matrix-object reading

Figure 2.1: Example of the picture selection task used in Tsimpli et al. (2004) for the item in (17a).

When presented with overt pronouns, the control groups for each language were expected to preferentially select matrix-object or new-referent interpretations (i.e., topic-shift readings). When presented with a null pronoun, both groups were expected to exhibit a bias for matrix-subject interpretations (i.e., non-topic-shift readings). For the two experimental groups, the authors predicted that their L2-English would influence

the discourse features that condition the interpretation of overt pronouns. Specifically, they predicted a bleaching of the [+TS] feature, resulting in indeterminacy in their interpretation. This is graphically sketched in Figure 2.2.

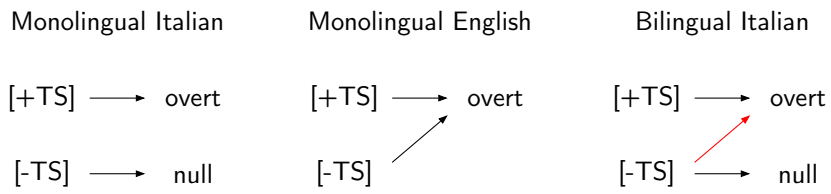

Figure 2.2: Sketch of the [±TS] specification of pronominal forms in monolingual Italian, monolingual English, and bilingual Italian.

For this task, however, the authors present data from the Italian groups but not the Greek ones. This appears to have been motivated by the fact that 'in some cases, significant effects were obtained only for one of the two languages' (Tsimpli et al. 2004: p. 267). To analyse their results, the authors modelled their four conditions (i.e., null and matrix-first, overt and matrix-first, null and embedded-first . . .) separately. Within each condition, they conducted ANOVAs over the selection rates of the three image types for a 2x3 design: *group* (control vs experimental) by *referent* (matrix-subject, matrix-object, new-referent).

For embedded-first items (17b) with null pronouns, the authors report a significant effect of *referent* which they interpret as a clear preference for matrix-subject readings in both L1-Italian groups (control: 85% of items; experimental: 87%). For those items, the effect of *group* and its interaction with *referent* were non-significant. For items in the same direction with overt pronouns, a significant *group* by *referent* interaction was observed. They report that this was driven by a significantly stronger preference for new-referent readings in the control group than in the experimental group (control: 63.6%; experimental: 39.9%). No differences were reported for the rates of matrix-subject or matrix-object readings.

For matrix-first items (17a), an interaction of *group* by *referent* was observed for items with both the null and overt pronouns. Within null items, they report that the experimental group opted for significantly more matrix-subject readings (control: 50.8%; experimental: 69.8%) as well as significantly fewer matrix-object readings than the control group (control: 44.1%; experimental: 24.9%). No difference was reported in the rates of new-referent readings. For items in the same direction with overt pronouns, they reported that their attrition effect was driven by the experimental group selecting significantly more matrix-subject readings that the control group (control: 7.6%; experimental: 21.2%). No differences were reported between matrix-object (control: 82%; experimental 72%) or new-referent readings (values $< 10\%$[6]).

The original authors interpret these results to support their hypothesis for two reasons. First, they observed an attrition effect in the interpretation of overt pronouns in both embedded- and matrix-first items. Second, no effect of attrition was observed in the embedded-first items with null pronouns. However, we should also note that are some inconsistencies with their bleaching account. First, the attrition effect for overt pronouns in the embedded-first items (17b) is not straightforwardly capturable in terms of a change in the [+TS] specification of overt pronouns. That trend was driven by fewer new-referent interpretations by the experimental group who instead selected numerically, but not significantly, more matrix-object and matrix-subject interpretations. Therefore, we cannot conclude that the experimental group admitted significantly more non-topic-shift interpretations than the control group in those items. Second, the increased matrix-subject interpretations for overt pronouns by the experimental group in the matrix-first items (17a) was modest (control: 7%; experimental 21%). As such, the experimental group is still exhibiting a (weakened) topic-shift bias in their interpretation of overt pronouns. This is not what would be expected if the feature had become underspecified. Third, the results from the matrix-first items (17a) with null pronouns cannot be

[6] These values are estimated from their Figure 7 as exact values were not provided.

easily captured under their feature bleaching account. Their experimental participants exhibited an increased non-topic-shift bias. To explain that issue, the authors instead speculate that the experimental group may have misanalysed the embedded CP as being nonfinite, the equivalent of (18), resulting in obligatory control by the matrix subject. We will return to this point after considering more data.

(18) The elderly woman greets the girl when cross**ing** the street.

In that same study, Tsimpli et al. (2004) also report the results from their Greek participants from the other half of the picture selection task. In that half, they investigated a related property of Italian-style null-subject languages, namely the syntactic optionality of pre- and postverbal subjects, and their resulting interpretive biases. To that end, an initial sentence always introduced a set of possible referents. The second sentence then always contained an indefinite subject DP that could be interpreted as old (i.e., one of the previously introduced referents) or new information (i.e., some unmentioned referent). In half of their items, the subject of the second sentence appeared preverbally and in the other half, it appeared postverbally. An example in both conditions is provided in (19). These were presented in the same manner as the pronominal items reported for the Italian participants and participants had to select between images corresponding to the two logically possible readings (old information, new information). For consistency with the other half of the picture selection task, the authors also included a third option which was invalid (in which there was no referent corresponding to the subject of the second sentence). An example of the three images for the item in (19) is provided in Figure 2.3.

(19) *I gitonisa mu ston trito orofo apektise dhidhima.* [Greek]
the neighbour my on.the third floor had twins

'My neighbour on the third flood had twins.

a. *Xtes vradhi ena moro ekleje.*

b. *Xtes vradhi ekleje ena moro.*
last night cried one baby

'Last night a baby cried.'

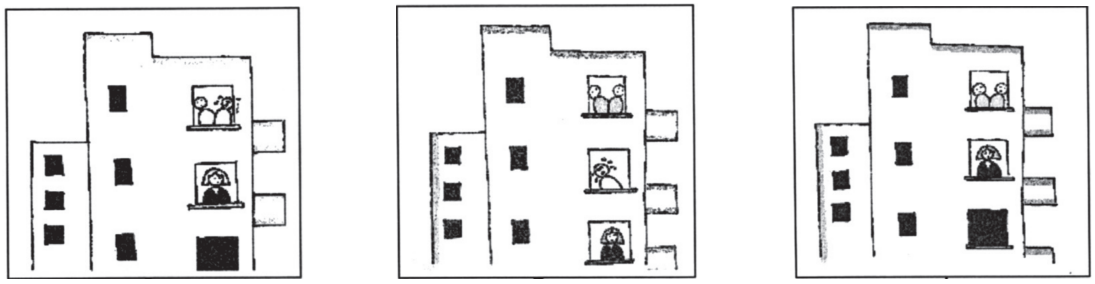

(a) old reading (b) new reading (c) invalid reading

Figure 2.3: Example of the picture selection task used in Tsimpli et al. (2004) for the item in (19).

For this portion of the task, the control group was expected to prefer 'old' readings for preverbal subjects and exhibit no clear bias for postverbal subjects as those subjects have been reported to be ambiguous between 'new' and 'old' readings in Greek. Under attrition, the experimental group was predicted to differ significantly in their interpretation of preverbal subjects only. In that case, they were expected to admit more 'new' readings under the influence of English which requires a preverbal subject regardless of information status. To analyse their data, the authors conducted ANOVAs over the selection rates for the relevant images for at least a 2x4

design: *group* (control vs experimental) by *referent* (new, old, both, and invalid). As it is unclear if two subject conditions (i.e., pre- and postverbal) were analysed separately or together, however, we refrain from interpreting their interaction terms and instead focus on the significant *post hoc*s for which it is clear which sections of the data were analysed.

For the preverbal items, the authors reported significant differences in their rates of 'new' and 'both' responses. Whereas the experimental group selected 'new' readings significantly less frequently than the control group (approximately[7] 10% of responses vs 28%), they selected 'both' readings significantly more frequently than the control group (approximately 42% vs 18%). No difference was reported for 'old' readings (approximately 44% vs 48%).

For the postverbal items, both groups exhibited indeterminacy, but in different ways. For the control group 'new,' 'old,' and 'both' were selected at similar rates (approximately 33%, 33% and 24% respectively). Conversely, the experimental group exhibited a clear bias for 'both' interpretations (approximately 52%) and opted for 'new' and 'old' readings around a fifth of the time (approximately 20% and 18%). In this condition, the two groups differed significantly in their rates of 'old' and 'both' readings.

Again, Tsimpli et al. (2004) interpret their results to support their hypothesis that attrition affects interpretable features only. They base this conclusion on the increased rate of 'both' readings for preverbal subjects which they take to indicate increased indeterminacy in the interpretation of preverbal subjects with regard to the feature [±Topic]. However, we should also note that their *post hoc*s indicated no difference in 'old' readings while also indicating that the experimental group selected significantly fewer – not more – 'new' (i.e. non-Topic) readings. Thus, although their interpretation is indeed possible and cannot be fully excluded based on the available data, it seems more reasonable to attribute the increase in 'both' responses to the decrease in pure non-Topic readings. To see this more clearly, consider what would have happened if the authors had only considered the overall response rates of 'new' and 'old,' rather than treating 'both' as a separate category. The control group would have shown a modest bias in line with the expected preference for 'old' readings (new: 46% of items; old: 66%[7])[8] whereas the experimental group would have shown a much more noticeable bias (new: 52%; old: 86%). In this particular case, the choice of how to present and analyse data affects which interpretations become more salient. Moreover, it is worth pointing out that we are not proposing a radical departure from the Tsimpli et al.'s (2004) own presentation of their data. It appears this is exactly what they have done for the pronominal items that they report for Italian. Recall, multiple responses were allowed and even encouraged, but no 'both/all' column is presented.

Applying the same logic to the results from postverbal items also affects our interpretation, albeit to a lesser extent. On the one hand, Tsimpli et al. (2004) interpret the experimental group's (i) significantly higher rates of 'both' responses and (ii) lower rates of 'old' responses to indicate increased indeterminacy. On the other, if we were to only consider the overall rates of 'new' and 'old' readings collapsing 'both,' we see that the two groups pattern together and appear to show no bias. This is what was predicted by Tsimpli et al. (2004). Overall, the control group selected 'new' readings in around 57% of items. Similarly, they selected 'old' readings around 57% of items. Looking at the experimental group the rates would be 72% and 70% respectively. Thus, it seems that the two groups do not differ in their interpretive biases (both showing indeterminacy) but the experimental group is overall more permissive (i.e., they accept more possible readings). We tentatively suggest that this might indicate that the experimental group is more sensitive to the interpretive ambiguity in their L1.

Integrating the Greek and Italian results then, it seems that only one of Tsimpli et al.'s (2004) attrition effects (matrix-first items with overt pronouns in Italian) provides clear evidence that attrition surfaces as the weakening of restrictions relating to discourse (i.e., interpretable) features. Conversely, if the reader is

[7] Exact values are not provided. As such, we report estimated values from their Figures 2 and 3.

[8] If we calculated the percentage over the total number of responses instead of over items such that the %s added to 100%, the values would trivially change but the overall pattern would not.

willing to accept our re-interpretations presented above, two of their attrition effects (matrix-first items with null pronouns in Italian and preverbal subjects in Greek) appear to indicate that attriters exhibited a more pronounced version of an interpretive bias already present in the L1 which cannot be ascribed to a relaxation of discourse features under influence of the L2. Finally, the remaining two effects do not appear to relate to (interpretable) discourse features.[9] First, there was the decrease in 'other' responses for items with overt pronouns in Italian which we should not attribute to a change in the [+TS] feature due to the lack of a significant difference in the rate of subject readings. Second, there was the increased acceptance rate for any possible interpretation of postverbal subjects in Greek. Interestingly, this is the only condition in which the control participants of the same L1 exhibit no clear bias. As such, we have suggested that this may relate to the fact that their L1 does not provide a bias (based on discourse feature or otherwise) to resolve the ambiguity.

A similarly mixed pattern can be observed in Gürel (2004). In that study, they were interested in how prolonged exposure to English might affect pronominal resolution in Turkish, another null-subject language.[10] In addition to null pronouns (ø), Turkish has two overt pronominal forms that may occupy the subject position: *o* ('s/he') and *kendisi* ('self'). Relevant for the present discussion, when these forms appear in contexts like in (20), *o* cannot be coreferential with the matrix subject. As such, it must be interpreted as disjoint from the referent of the matrix subject (i.e., someone other than Burak). This contrasts with pronominal subjects in English (21) which admit either interpretation. The same asymmetry does not hold for *kendisi* or ø, which may corefer with the subject or receive a disjoint interpretation.

(20) *Burak$_i$ [(o-nun$_{*i/j}$ / kendi-si-nin$_{i/j}$ / ø$_{i/j}$) zeki ol-duğ-u]-nu düşün-üyor.* [Turkish]
Burak (s/he-GEN / self-3.SG-GEN / ø) intelligent be-NOM-3.SG.POS-ACC think-PRG

'Burak$_i$ thinks that [he$_{*i/j}$/self$_{i/j}$/ø$_{i/j}$[11] is intelligent].'

(21) John$_i$ believes [he$_{i/j}$ is intelligent].

For their study, Gürel recruited an experimental group of native Turkish speakers (N = 24) who had immigrated to Canada or the US after the age of 16. At the time of testing, they had a mean age of 47 years and had resided in their L2-speaking countries for a minimum of 10 years (mean = 21.5 years). L2 proficiency was not assessed. Rather the author assumed the participants to be highly proficient in their L2-English given their professional and educational backgrounds. Moreover, no explicit information was given regarding the estimated usage of the L1 and L2. For control participants, they also recruited a group of native Turkish speakers still residing in Turkey (N = 30) who had a mean age of 40. Notably, this group did not consist of idealised monolinguals as they spoke some English and were deemed competent enough in their L2-English to take part in the story task described below.

Participants were presented three tasks. In the written interpretation task, participants were provided with a series of sentences as in (22). Within those items, the author manipulated the pronominal form (*o*, *kendisi*, ø) and whether the potential antecedent was referential or quantified (e.g., Burak vs everyone). After each sentence, participants were asked to indicate their interpretation of the embedded subject selecting from three options: coreferential with the matrix subject (e.g., Burak), disjoint reference (i.e., someone else), or both.

[9] In the original paper, the authors also report a third effect which does not appear to be directly linked to discourse features. Namely, in their 'headlines' task, the Greek experimental group exhibited an unexpected definiteness effect in their production of pre- and postverbal subjects. It is unclear how to integrate that into the Greek results from the picture selection task, as the production task uniformly elicited production in all-new-information contexts.

[10] Although Turkish allows elision of object DPs, see Şener & Takahashi (2010) for a discussion on the licensing of null subjects via agreement rather than elision, akin to Spanish or Italian.

[11] Following Gürel (2004), I will copy the Turkish indices into the English translation for clarity regarding the readings we are interested in.

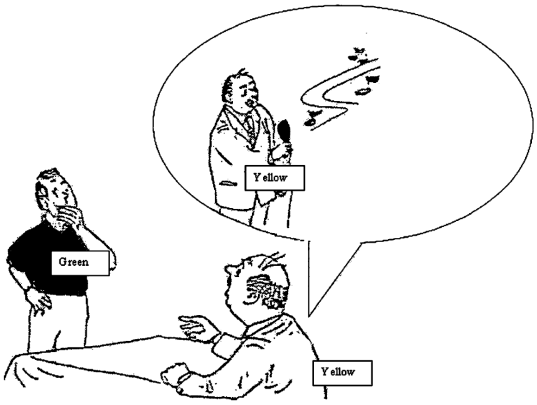

Figure 2.4: An example image for the corresponding sentence in (24).[12]

(22) *Burak$_i$ [o-nun$_{*i/j}$ sinema-ya gid-eceğ-i]-ni söyle-di* [Turkish]
Burak s/he-GEN cinema-DAT go-NOM-3.SG.POS-ACC say-PST

'Burak said [s/he$_{*i/j}$ would go to the movies].'

Similar items were used in the story-based truth-value judgement task. However, in that case, participants were presented a short story in English prior to each item. This was done to provide a context in order to force particular interpretations of the pronominal elements. Against these story contexts, participants were asked to judge whether the following Turkish sentence was true or false. An example of an English story (forcing disjoint reference) with a corresponding Turkish sentence is provided in (23).

(23) Mary and John went to a restaurant. Mary ordered seafood and John ordered a pizza. The bill came to 50 dollars. John complained that the bill was high but Mary didn't agree.

a. *Mary$_i$ [o-nun$_{*i/j}$ restoran-ı pahalı bul-duğ-u]-nu söyle-di.* [Turkish]
Mary s/he-GEN restaurant-ACC expensive find-NOM-3.SG.POS-ACC say-PST

'Mary said s/he$_{*i/j}$ found the restaurant expensive.'

In the third task, a picture-based truth-value judgement task, participants were presented a single image per item to force particular interpretations of the pronoun. They then listened to a pre-recorded sentence in Turkish and had to indicate if the sentence was true based on the picture. Only referential potential antecedents were considered due to difficulties in drawing quantified antecedents. An example of an image for the sentence in (24) is provided in Figure 2.4.

(24) *Ahmet Sarı$_i$ [o-nun$_{*i/j}$ iyi şarkı söyle-diğ-i]-ni söyle-di* [Turkish]
Ahmet Sari s/he-GEN well song tell-NOM-3.SG.POS-ACC say-PST

'Ahmet Sari$_i$ said that s/he$_{*i/j}$ sings well.'

Results from the written interpretation task indicated that both groups almost always selected the disjoint reading and almost never interpreted *o* as coreferential with the matrix subject. Statistical analysis indicated that the only difference between the two was that the experimental group selected subject interpretation (6% of trials) marginally more frequently than the control group (2%) in quantified items. No difference was found for referential antecedents (1% for both groups). Moreover, no statistically significant differences were found in the rates of 'both' (i.e., subject or disjoint) readings.

[12] It appears that the image was originally in colour to help participants identify Ahmet (who was in yellow).

As for the *kendisi* with referential and quantified antecedents, the control group preferred 'both' interpretations (64% and 68%) and never selected disjoint readings. The experimental group similarly dispreferred disjoint readings (3% and 7%). However, they rarely gave 'both' responses and instead preferred subject interpretations (81% and 63%). A similar, but less pronounced, shift away from 'both' and toward subject readings by the experimental group was also attested for null pronouns. For both pronominal types, these shifts toward more subject readings were significant.

For the story-based truth-value judgement task, both groups preferred disjoint reading for *o* and subject readings for *kendisi* and ø. The only significant difference between the groups surfaced with *o*. While the control group almost never accepted subject readings (referential: 4%; quantified: 3%), the experimental group did so at a higher rate for both referential and quantified antecedents (30% and 22%). Similarly, in the picture-based task (which only tested referential antecedents), the control group categorically rejected bound readings whereas the experimental group accepted such readings in a subset of trials (21%). However, in that task, significant differences were also observed within *kendisi* and null pronominal items. Although both groups preferred subject readings, this bias was stronger in the experimental group (*kendisi*: 85% vs 97%; ø: 78% vs 92%). This is convergent with the increased subject bias observed in the written interpretation task.

Gürel (2004) interprets her results to indicate three things. First, the distinction between the *o* and the other two pronouns is maintained under attrition. Second, attriters begin to treat *o* more like 'he' under pressure from English, potentially due to the overlap and competition between the linguistic systems. Third, they take the increased bias toward subject readings for *kendisi* and ø to indicate that the experimental group exhibited a reduced awareness of the ambiguity for those pronouns.

Following up on that original study, Gürel & Yılmaz (2011) re-ran the written interpretation task from Gürel (2004) with potential attriters and HSs of Turkish living in the Netherlands. Dutch, like English, is a non-null-subject language in which overt pronouns in the relevant contexts may be interpreted as coreferential or disjoint with the matrix subject. Compare (25) with (20) and (21).

(25) *Ashley$_i$ gelooft dat [zij$_{i/j}$ charmant is].* [Dutch]
Ashley believes that she charming is

'Ashley$_i$ believes that she$_{i/j}$ is charming.'

Focusing on Gürel & Yılmaz's (2011) potential attriters (N = 19), participants had emigrated to the Netherlands after the age of 15. At the time of testing, they had been living there for an average of 19.89 years and had an average age of 42.05 years. Most (N = 15) were deemed to be highly proficient in the L2-Dutch[13] which they had acquired after emigration. For a control group, the authors re-used the data from the control group in Gürel (2004). Moreover, their analysis also included the original North American experimental group for comparison. However, given that Gürel (2004) observed minimal differences between referential and quantified antecedents, the distinction was collapsed in Gürel & Yılmaz (2011).

Results indicated that the new European experimental group did not statistically differ from the control group with regard to their interpretation of *o*. That is to say, the experimental group overwhelmingly selected disjoint interpretations (86% of trials) and almost never selected subject interpretations (3%). For *kendisi* and ø items, the experimental group again patterned with the control group in that they almost never selected disjoint readings (3% for both pronouns). However, whereas the control group preferentially selected 'both' readings (64% and 84% respectively) for the relevant items, the European experimental group selected those interpretations at a significantly lower rate (10% and 32%). They instead overwhelmingly selected more subject readings (87% and 74%). As such, the European participants pattern closely with the North American ones from Gürel (2004). Moreover, when the two experimental groups were directly compared *post hoc*, no significant differences were observed. As such, Gürel & Yılmaz (2011) were able to replicate the findings in

[13] Based on the Dutch Language Proficiency Exam NT2 Programma I & II

Gürel (2004). We take this to indicate that the shift toward more subject interpretations of ø and *kendisi* is a general pattern in Turkish under attrition.[14]

Integrating the Turkish findings with those from Tsimpli et al. (2004), there are two patterns worth highlighting. Starting with the overt pronoun *o*, the results from Gürel's (2004) two truth-value judgement tasks indicate that *o* in Turkish is susceptible to attrition similar to overt pronouns in Italian, with matrix subject readings becoming more prevalent. However, to this, we should append two caveats. First, the effect in Turkish (story-based: referential: 4% vs 30%; quantified 3% vs 22%; image-based: referential: 0% vs 21%) is similar in magnitude to the effect observed in matrix-first items with overt pronouns in Italian (7.6% vs 21.2%). Clearly in neither case have the attriting participants lost the L1 bias, despite its weakening. Second, it is worth noting that the same participants did not exhibit an attrition effect with *o* in Gürel's (2004) written interpretation task (it was marginal, and only in half of the items). Moreover, neither did the participants in Gürel & Yılmaz's (2011) replication thereof. We might suggest two possible explanations for the asymmetry. One option would be to speculate that task differences played a role. In the written interpretation task, participants were asked to select their preferred interpretation from a set of options (subject, other, both). Conversely in the truth-value judgement tasks, participants had to actively reject sentences that did not conform to their interpretive biases. This is arguably a cognitively more demanding task which may have played a role in the results. Alternatively, it may again relate to the coding of the data. Namely, the written interpretation task allowed for 'both' responses whereas the truth-value judgement task trivially could not. In the written task, not only did the experimental groups provide numerically more matrix-subject responses than the control group, they also provided numerically more 'both' responses (control: 5%; North American: 15%; European: 11%). It is possible that different results would have been found had the authors collapsed 'both' into their constituent matrix-subject and disjoint readings.

A second pattern worth highlighting from the Turkish data relates to the null pronoun and *kendisi*. Starting with the written task in Gürel (2004), results indicated that overall the experimental group preferentially selected matrix-subject readings whereas the control group preferentially selected 'both' readings. In line with Gürel (2004) we might interpret this as reduced awareness of the ambiguity of such forms in the L1. However, if we integrate the results from their two truth-value judgement tasks we might come to an alternative interpretation. Namely, in those tasks where participants were forced to accept or reject a particular reading, participants did not exhibit indeterminacy. Rather, all participants (experimental and control) preferentially interpreted the null pronoun and *kendisi* as referring to the matrix subject (all values $\geq$ 74%). As such, it seems reasonable to suggest that, although not categorically, Turkish speakers exhibit a matrix-subject bias in the interpretation of the relevant pronouns. Integrating this with the written task results from Gürel (2004) and Gürel & Yılmaz (2011) then, we might suggest that the experimental participants are not less aware of the ambiguity, but instead more sensitive to their L1's matrix subject bias for the relevant forms. Under that interpretation, the results for the null pronoun and *kendisi* in Turkish are convergent with Tsimpli et al.'s (2004) results for the matrix-first null pronominal items in Italian and the preverbal items in Greek. This suggests that the tendency toward increased interpretive biases for certain structures is not spurious and should be accounted for.

A similar pattern was also observed in a pilot study conducted by Martín-Villena (2023) to explore how different types of temporal adverbial clauses affect the PAH. For this, the author recruited native (Iberian) Spanish speakers (N = 131) who had grown up monolingually and were still residing in Spain (Martín-Villena, p.c.). At the time of testing, participants had a mean age of 32.2 years and reported a mean daily Spanish usage of 86.5%. Within these participants however, 76 considered themselves 'highly proficient' in their

[14] We are aware of Kim, Montrul & Yoon's (2009a) work on *caki* (for the present purposes, analogous to *o* in Turkish) in the Korean of potential attriters. However, we do not take that to detract from the relevant Turkish pattern for two reasons. First, their group of potential attriters was noticeably smaller (N = 10). Second, unlike Kim, Montrul & Yoon's (2009b) study with HSs, they do not appear to have tested *caki-casin* (for the present purposes analogous to *kendisi* in Turkish). As such, they do not bear on the increased subject bias for the relevant forms.

L2-English, while 55 considered themselves as 'not [being] proficient enough in English.'

The experimental method and items were adapted from the Greek version in Tsimpli et al. (2004). As discussed above, those items always contained matrix CP with the two possible antecedents. In a slight departure from the original task, the matrix CP was always followed by the embedded CP which contained the ambiguous pronoun. In half the items, the pronoun was overt. In the other half, it was null. To explore the role of different adverbials, half the items contained *mientras* ('while') and half *cuando* ('when,' 26). These items were presented one at a time on the participants' screen along with three images corresponding to subject, object, and 'other' interpretations of the ambiguous pronoun. Participants were instructed to select their preferred interpretation (i.e., only one unlike in Tsimpli et al. 2004), after which the experiment progressed to the next item. An example of the subject and object images for the item in (26) is presented in Figure 2.5.[15]

(26) a. *La anciana saludó a la mujer cuando ella cruzaba la calle.* [Spanish]
the elderly.woman greeted DOM the woman when she crossed the street
'The elderly woman greeted the woman when she crossed the street.'

b. *El padre saludó al hijo mientras él montaba en bicicleta.* [Spanish]
the father greeted DOM.the son while he mounted on bike
'The father greeted the son while he was riding a bike.'

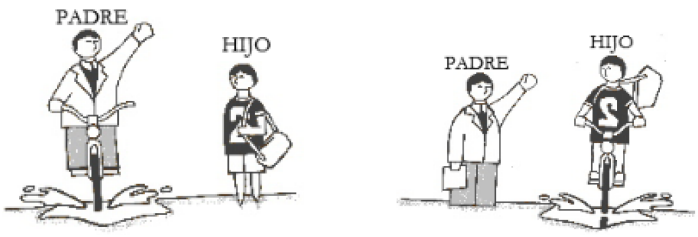

(a) SpecIP reading (b) Non-SpecIP reading

Figure 2.5: Example of the images used for the picture selection task in Martín-Villena (2023) for the item in (26b).

In their analysis of the data (a generalised logistic regression), the original author only considered subject and object responses due to the paucity of 'other' responses (3%). In their analysis, they entered 4 predictor variables: pronoun (null vs overt), conjunction (*cuando* 'when' vs *mientras* 'while'), L2 proficiency (high vs not-high) and age (continuous). Statistical modelling indicated a significant effect of *pronoun*, surfacing as fewer SpecIP interpretations for overt pronouns. There was also a significant effect of *conjunction* indicating more SpecIP interpretations for *mientras* ('while'). They also report three significant two-way interactions with *pronoun*.[16] *Post hoc* pairwise comparisons indicated that their *pronoun* by *conjunction* interaction was driven by an increased SpecIP bias for null *mientras* ('while') items. They also report that the *age* by *pronoun* interaction was driven by more SpecIP interpretations by the older participants within the overt conditions. Most relevant to the present discussion, however, they also reported a significant *L2-level* by *pronoun* interaction. *Post hoc*s indicated that this was driven by fewer SpecIP interpretations for the highly proficient L2ers in the overt conditions.[17] That is to say, participants who deemed themselves more proficient

[15] The image corresponding to an 'other' reading was not provided.

[16] Their model did not include any potential interaction terms without *pronoun* (i.e., *conjunction* by *age*). Nor did it include any higher order interactions (i.e., a three-way interaction between *pronoun*, *conjunction*, and *L2-level*).

[17] Raw values are not provided. From their Figure 38, it appears that values predicted by the model are approximately: non-high L2ers: 19%, high L2ers: 10%. It is unclear if those predicted values are for the reference level for conjunction (*cuando*) or if the conjunction distinction was collapsed. We take the latter to be more likely.

in their L2 exhibited a stronger – not weaker – non-SpecIP bias in their interpretation of overt pronouns. This is unexpected under a feature-based account but fits with the other pattern observed in the Greek, Italian, and Turkish results above. Namely, it suggests that increased proficiency in an L2 may lead to increased interpretive biases for both null and overt pronouns which is not attributable to the transfer from the L2 (either at the level of representations or processing). Of course, one might reasonably object that even though these 'highly proficient' L2ers are bilinguals, they still resided in their L1 community and as such may not be representative of prototypical attriting populations in migration contexts. To that end, we now turn to Martín-Villena's (2023) main study in which they compared potential attriters living abroad with those still in their L1 community but received extensive L2 contact.

With the *pronoun* by *conjunction* interaction from the pilot in mind, Martín-Villena (2023) then conducted an almost identical[18] picture selection task to explore potential attrition effects using three well-defined groups of native (Iberian) Spanish speakers: a control group, an immersed experimental group, and an instructed experimental group. Starting with the control group, participants consisted of 'functionally monolingual' speakers (N = 33). Those participants had never lived abroad and had minimal contact with English. This was restricted to state-mandated instructed settings in primary and secondary education with an average of 11.6 years of lessons (mean age of onset = 5.82 years). Their L2 proficiency was assessed via the Oxford Quick Placement Test (OQPT) with participants receiving a mean score of 22.2/60 (max = 29). This indicates that their L2 proficiency corresponds to an A1 or A2 within the Common European Framework of Reference for Languages (CEFR). Moreover, to assess language dominance, participants completed the Bilingual Language Profile (BLP) questionnaire developed by Birdsong, Gertken & Amengual (2012) which asks questions about language history, use, proficiency, and attitudes. From this, an index of dominance (from -218 to 218) may be obtained where positive values indicate that a participant is more L1 dominant and negative values indicate that they are more L2 dominant. For the control group, the mean BLP score was 143 indicating they are highly L1 dominant. At the time of testing, these participants had a mean age of 21.2 years.

For their immersed bilingual group, the author recruited participants (N = 94) who were living in the United Kingdom or the Republic of Ireland at the time of testing. For this group, their minimum age of immigration was 15 years. However for the vast majority of participants (N = 90) this occurred after their 18th birthday. Unlike previous work (e.g., Chamorro et al. 2015a, Gürel 2004, Gürel & Yılmaz 2011, Tsimpli et al. 2004) no minimum length of residence in the L2 was imposed on this group. Rather, participants had a range of residencies: 1 (N = 16), 2 (N = 16), 3 (N = 20), 4 (N = 17) or 5+ years (N = 25). Like the control group, these participants had received English instruction as part of their education (mean = 13.9 years, mean age of onset = 6.2 years). English proficiency was again assessed via the OQPT. Their mean score was 52.7/60 indicating that they likely have a C1 or C2 level within the CEFR. Moreover, participants reported that their mean daily usage of English was 64.5%. They also completed the BLP and received a mean score of 44.7 indicating that they are still L1 dominant although to a lesser degree than the control group. At the time of testing the immersed group had an average age of 26.6 years.

The instructed experimental group consisted of participants (N = 80) still residing in Spain. As with the other two groups, all these participants had been raised monolingual but had received instruction in their L2 English for an average of 15 years (mean age of onset = 5.21 years). Moreover, at the time of testing, participants in this group were enrolled in an English Studies degree at Spanish universities with roughly 20 participants in each year of their 4-year course. Content for these degrees was primarily delivered in English. Participants were screened to ensure regular attendance to lectures (approximately 20 hours per week) and that none had spent a semester or more abroad in an English-speaking country. L2 proficiency was assessed via the OQPT and on average participants scored 52.4/60 indicating that they have a C1-C2 level within the CEFR. This is comparable with the author's immersed group (52.7/60). However, the instructed bilinguals

[18] Given the trivial rate for 'other' responses in the pilot study, this option was removed from the main picture selection task.

noticeably differed in their reported daily English use (25.5% vs 64.5% for the immersed group). As for the BLP, participants had a mean score of 63.33, indicating that they are still L1 dominant and somewhere between the control and immersed groups. At the time of testing, the instructed group had a mean age of 20.4 years.

For their analysis (a generalised logistic regression) of the pooled data for the three groups, the model included effects of *pronoun* (null vs overt), *conjunction* (*cuando* 'when' vs *mientras* 'while'), and *group* (control, instructed, immersed) as well as the scaled continuous predictors *BLP score* and *working memory* (as tested by a sentence-reading-span task). The model additionally included the following potential interactions: *pronoun* by *conjunction*, *pronoun* by *group*, *pronoun* by *conjunction* by *group*, *pronoun* by *BLP score*, and *pronoun* by *working memory score*.

Abstracting away from some of their less relevant results (e.g., greater rates of subject responses as working memory scores increased), the authors reported a significant *pronoun* by *conjunction* interaction for all of their groups. In line with their pilot study, this was interpreted to indicate a greater rate of subject responses for null subjects when the conjunction was *mientras*. Also in line with their pilot study, the author reports that *post hoc* tests[19,20] indicated that (i) both experimental groups selected subject readings for overt pronouns significantly less frequently than the control group (control: *mientras*: 30%, *cuando*: 23%; instructed: *mientras*: 17%, *cuando*: 15%; immersed: *mientras*: 19%, *cuando*: 20%), and (ii) the two experimental groups did not differ from each other in this regard. The author does not report whether the effect of *group* was significant within the null items.

As for the interaction of *pronoun* with their continuous variables *BLP score* and *working memory score*, both terms were found to be significant. Starting with the *pronoun* by *working memory score* interaction, they observed that increased working memory was associated with increased subject responses for null pronouns only. Turning then to the *pronoun* by *BLP-dominance* interaction, the author observed that this surfaced as fewer subject interpretations for overt pronouns as BLP scores (i.e., dominance in Spanish) increased. On the surface, this seems to converge with previous work that has reported increased subject interpretations/reduced PAH biases for overt pronouns under attrition. However, this conflicts with two of the same author's previous findings. First, the results from the pilot study indicated fewer subject interpretations for overt pronouns in the 'highly proficient' L2ers. Second, *post hoc*s from the experiment presently under discussion indicated fewer subject interpretations for overt pronouns in both experimental groups compared to the functional monolinguals. Recall also that BLP scores were noticeably different between the three groups with the control group's mean score being higher than the other two groups' (control: 143; instructed: 63.3; immersed: 44.7). Martín-Villena (2023) interprets this conflict as evidence that we should move beyond dichotomous groups and instead investigate attrition with continuous measures of dominance. However, the interaction with *BLP* does not negate the interaction with *group*. Moreover, if dominance (as measured by an index like the BLP) is the factor that we should be interested in when exploring attrition rather than the traditional factor of group, one could make the prediction that the interaction with *BLP* should survive in a simplified model which does not include *group*. This is not the case (Martín-Villena, p.c.); the significance of the interaction with *BLP* depends on the inclusion of the interaction with *group*, but not *vice versa*. That observation in turn raises the question of what relevant aspect of the participants is the BLP tapping given that BLP is trivially related to group membership. This is compounded by the fact that exploratory modelling indicated that none of the four subcomponents of the BLP (e.g., current language use) revealed significant interactions with *condition* on their own.

Nonetheless, as Martín-Villena's (2023) interpretation of the *pronoun* by *BLP score* interaction differs from ours, they also present analyses over the immersed and instructed bilinguals individually to further explore

[19] It is unclear how they conducted their *post hocs* and whether these were run over the reference level for conjunction (i.e., *cuando*), or if the distinction between the two conjunctions was collapsed. We take the latter to be more likely.

[20] Although not reported, significant *group* by *pronoun* interactions were observed between the two experimental groups and the control group indicating that these *post hoc*s were warranted (Martín-Villena, p.c.).

that interaction. For both groups, data was analysed with a logistic regression similar to the one above with two changes. First, they trivially do not consider *group* as the groups are analysed separately. Second, they additionally consider *length of residence (LoR)/length of instruction (LoI)* as new potential interaction terms. Most germane to the present thesis, analysis of the instructed group indicated a significant interaction of *pronoun* and *BLP score*. Based on visual inspection, they interpret this to indicate that higher L1 dominance was associated with fewer subject interpretations for overt pronouns. For the immersed group, the simple two-way interaction of *pronoun* and *BLP score* was not significant, but their three-way interaction with *LoR* was. Again based on visual inspection, they attribute this to an effect of *BLP score* within the longer immersed bilinguals with the more L1 dominant bilinguals selecting fewer subject interpretations for overt pronouns. Thus, these results seem to be in line with the results from the model that contained data from all three groups. As such, the author interprets their data to fit the predictions of the IH. Namely, attrition affects overt pronouns only, with more subject interpretations for overt pronouns as participants shift toward L2 dominance.

Given their own modelling, however, we should be extremely cautious of the latter half of that author's conclusion. First, as noted above the author's own group level *post hoc* tests indicate that the immersed and instructed bilinguals exhibit more extreme non-subject biases for overt pronouns than the control group do despite the control group exhibiting a lower BLP. Second, although the Martín-Villena (2023) takes the BLP as an index of dominance, it is not entirely clear what the BLP is indexing in those participants. Without a re-analysis of their data (particularly the subcomponents of the BLP),[21] their published findings could be used to argue for opposing conclusions. On the one hand, the authors interpret their BLP interactions to indicate a relaxation of the non-SpecIP bias for overt pronouns in line with Tsimpli et al.'s (2004) results for overt pronouns in Italian, and Gürel (2004) and Gürel & Yılmaz's (2011) results for *o* in Turkish. On the other hand, the group-level *post hoc*s from their pilot and main study indicated that participants who are more proficient in their L2 exhibit an increased sensitivity to their L1's non-SpecIP bias for overt pronouns in line with (i) Tsimpli et al.'s (2004) results for null pronouns in Italian (ii) the same authors' results for preverbal subjects in Greek as well as (iii) Gürel (2004) and Gürel & Yılmaz's (2011) results for ø and *kendisi* in Turkish.

2.3.3.2 Attrition and (surface) overlap with the L2

Above we have argued for two contrasting patterns. On the one hand, attrition is sometimes reported to result in the overextension of overt subjects to non-topic shift contexts. On the other hand, attrition appears to also lead to increased interpretive biases in some cases, particularly in the case of the non-topic-shift bias for null pronouns. This raises the question of why the interpretive biases for different aspects of null-subject languages should be differently affected. One potential avenue to resolve this tension is to follow Sorace (2011) in proposing that overt pronouns are differently affected because they serve as a 'processing default' (although they are not the linguistic default). Recall from section 2.2 that monolingual speakers of Italian become noticeably more accepting of non-topic-shift readings for overt pronouns in single-referent contexts than in multiple-referent contexts (Carminati 2002). Moreover, the same speakers also appear to be more rigid in their interpretative preference for null subjects than overt ones in general. As such, Sorace (2011, 2016) suggested that the overextension of overt pronouns by bilingual speakers is not qualitatively different from that in monolingual speakers. Rather, under the assumption that processing pronominal reference in real-time is effortful, Italian speakers (in general) sometimes overextend overt pronouns to alleviate processing load when ambiguity is not at stake or when they experience (temporary) difficulties. Under such a view, the difference between bilinguals and monolinguals is quantitative, in that the former is more permissive of PAH violations, potentially due to the need to expend cognitive resources to maintain language control. An

[21] At the time of writing, their data was unavailable.

alternative potential factor that we should also consider is the overlap between the L1 and L2.[22] The first argument explicitly along these lines that we are aware of comes from Gürel (2004). Recall that under the influence of L2 English, that author observed that the overt pronoun *o* in Turkish is overextended whereas we have argued that the interpretive biases for the null pronoun and *kendisi* were clearer in the experimental group. To account for the overextension of *o*, Gürel (2004) makes use of Paradis' (1993) Activation Threshold Hypothesis (ATH). The ATH proposes that there is a correlation between how frequently a linguistic element is used and its activation (i.e., accessibility) for a speaker. Essentially, the less a given form is used, the higher its threshold for activation (i.e., it is less available) while the more frequently that a given form is used the lower its threshold (i.e., it is more available). Moreover, the ATH predicts attrition to occur when a form in the L1 has a 'competing' form in the L2 with a lower activation threshold (Paradis 2007). Relevant to the present discussion, this predicts that pronominal forms would be differentially affected by attrition. In Gürel's (2004) case, they assume that the overt pronoun *o* in Turkish competes with overt English pronouns such as *he*. Conversely, they assume that null pronouns and *kendisi* do not (compare 27 and 20). In this way, they can derive why only the interpretive restriction for *o* is relaxed under attrition.

(27) *Burak thinks that (self / ø) is intelligent.

Extending this idea to L1-Italian L2-English bilinguals, there is some initial evidence that attrition brings about an overextension of postverbal subjects only where there is (surface level) overlap between the L1 and L2, resulting in possible competition. Relevant to this discussion, it is usually assumed that the unmarked word order for Italian is subject-verb-object (SVO). However, there are several well-known contexts in which Italian also allows – or even requires – subjects to appear postverbally: (i) the single argument of an unaccusative verb may remain *in situ* in certain circumstances (with V-to-T movement resulting in a postverbal subject), (ii) the subject of any verb type may occupy the vP-periphery when it bears certain discourse features (e.g., narrow new-information focus), and (iii) certain wh-operators trigger obligatory subject inversion or dislocation in both matrix and embedded environments. This final case is exemplified in (28) with an unergative verb considered in an out-of-the-blue context to abstract away from the other two routes to postverbal subjects.

(28) a. **Dove Gianni dorme?* [Italian]

b. *Dove dorme Gianni?*
where sleeps Gianni

'Where does Gianni sleep?'

As discussed by Rizzi (2001, 2006), there are exceptions to this general ban on preverbal subjects in wh-questions, particularly for *perché* ('why,' 29), *come mai* ('how come'), and to a lesser extent for lexically restricted, D-linked[23] wh-phrases such as *chi di loro* ('who of them').

(29) a. *Perché Gianni dorme?* [Italian]

b. *Perché dorme Gianni?*
why sleeps Gianni

'Why is Gianni sleeping.'

[22] These two avenues are not mutually exclusive. However, should one wish for a single account on grounds of parsimony, this would require arguing that any potential overextension of overt pronouns in attriters whose L1 and L2 are both null-subject languages is the result of microvariation in the observance of the PAH as suggested by Filiaci et al. (2014).

[23] Following Pesetsky (1987), wh-phrases are said to be discourse/D-linked when they naturally prompt selection from a set of referents that have already been introduced to the discourse.

To experimentally investigate preferences for pre-/postverbal subjects in different types of wh-questions, Bocci & Pozzan (2014) conducted a forced-choice task with native speakers of Italian (N = 12). In that task, participants were presented a series of short exchanges between two speakers (A and B) as exemplified in (30). They were then asked to select between the two possible B alternatives. Stimuli only contained unergative verbs to avoid postverbal subjects via the internal argument position. Moreover, the preceding contexts (A) were designed not to induce focus or topicality interpretations to avoid postverbal subjects via the vP-periphery. Within the items, the authors manipulated the question type (*dove* – 'where,' *perché* – 'why,' and polar) as well as clause type (matrix or embedded question). Results indicated the expected preference for postverbal subjects for *dove* ('where') questions (values > 95%). This surfaced regardless of clause type. For *perché* ('why') and polar questions, however, participants exhibited a clear preference for preverbal subjects in both matrix and embedded contexts (all values > 70%).

(30) A. *Questa muscia è terribile! Non mi viene proprio voglia di andare in pista.* [Italian]
this music is terrible not to.me comes really desire of go on dance.floor

'This music is terrible! I really don't feel like dancing.'

B. *Hai proprio ragione . . . Non so perché Marta balla.*

B′. *Hai proprio ragione . . . Non so perché balla Marta.*
Have.2.SG really reason not know.1.SG why dances Marta

'You are totally right. I don't know why Marta dances.'

The authors also report convergent results from a naturalness judgement task with a different set of native Italian speakers (N = 108). That task used the same items as above and asked participants to rate how natural the various responses sounded. Results indicated that postverbal subjects were rated as significantly more acceptable than preverbal ones in the *dove* ('where') items. In the *perché* ('why) and polar questions, however, the opposite trend was observed.

More germane to the present thesis, the authors also reported a third experiment for which they recruited native Italian speakers (N = 12) who had been living in the US or UK for a minimum of 2 years. Speakers potentially undergoing attrition were selected using a questionnaire that asked about (i) difficulties with lexical access, (ii) using structures that might sound odd in Italian, and (iii) directly translating expressions from English to Italian. Of the final 12 participants, all but one indicated occasional difficulties with the above. Participants were then presented the same forced-choice task as the native Italian speakers still living in Italy. Results indicated no group difference in embedded questions; both groups preferentially selected postverbal subjects after *dove* ('where') questions (values > 95%) and preverbal subjects after *perché* ('why') and polar questions (values > 88%). In the matrix questions however, despite both groups categorically selecting postverbal subjects after *dove* ('where') questions (values > 95%), the two groups appeared to differ with regard to *perché* ('why') and polar questions; whereas the native Italians still living in Italy preferentially selected preverbal subjects (approximate[24] values: *perché*: 71%; polar: 86%), the experimental group selected preverbal subject noticeably less frequently (approximate values: *perché*: 40%; polar: 65%). When responses were statistically modelled (logistic regression), the three-way interaction of *group*, *context*, *question type* was significant. In the absence of *post hoc*s, this was interpreted as being driven by the difference between the two groups within the non-*dove* questions in the embedded context.

To account for the group difference in embedded contexts, the authors speculate that this may be related to the pattern of T-to-C movement in the L2. Namely in English, movement of the tensed element (T) to the left periphery (CP) is obligatory in matrix questions (of the relevant types) but blocked in embedded contexts (31). Thus, they suggest that conflict between the default word order in their L1 and the obligatory word

[24] Exact values are not provided. As such these are estimated from Figure [3], Study 3.

order in their L2 led to the observed shift in the potential attriters. However, this effect is restricted to those cases in which the L1 exhibits (putatively) syntactic optionality conditioned by discourse (i.e., interpretable) features. Where preverbal subjects are ungrammatical in the L1 but are obligatory in the L2 (i.e., embedded *dove* questions) the potential attriters exhibited an overwhelming preference for postverbal subjects ($>$ 95%) just like control participants.

(31) a. Why/where is Marta dancing?

b. *Why/where Marta is dancing?

c. *I don't know why/where is Marta dancing.

d. I don't know why/where Marta is dancing.

To follow up on Bocci & Pozzan's (2014) hypothesis that the optionality in subject position with *perché* ('why') questions as in (29) is conditioned by discourse (i.e., interpretable) features, Bianchi, Bocci & Cruschina (2017) conducted a forced-choice task similar to that in Bocci & Pozzan (2014) with native Italian speakers (N = 65). For this task, however, items were always preceded by either a neutral or focus favouring context as in (32). Participants were then asked to select between a pair of *perché* ('why') question responses with unergative verbs that differed only in the position of the subject (pre- or postverbal) akin to those in (29). The experiment also included *dove* ('where') and *come* ('how') question pairs in the neutral context condition for comparison.

(32) a. Neutral context

A causa di un problema tecnico hanno dovuto spostare la prova generale e le
for cause of a problem technical have.3.PL needed move the trial general and the
aule per le prove individuali sono state riassegnate, per ciu Giulia chiede al
rooms for the trials individual are been reassigned for which Giulia asks to.the
direttore: [Italian]
director

'Due to a technical problem, the dress rehearsal had to be postponed and the rooms for the individual rehearsals were reallocated, so Giulia asks the director:'

b. Focus favouring context

Giulia non sa che hanno cambiato il primo ballerino per il pas á deux e chiede
Giulia not knows that have.3.PL changed the first dancer for the *pas á deux* and asks
stupita:
surprised

'Giulia doesn't know that the lead dancer for the *pas á deux* was changed and ask surprised:'

As in Bocci & Pozzan (2014), results indicated a categorical preference for postverbal subjects in *dove* ('where') questions (98.2%) in neutral contexts. For *perché* questions in neutral contexts instead, participants only selected postverbal subjects in 37.4% of trials. That value is in line with what was reported by Bocci & Pozzan (2014) but clearly contrasts with the *perché* questions in focus favouring contexts. There participants preferentially selected postverbal subjects (66.5%). This suggests that in *perché* questions, subject position is conditioned by discourse features, and it may move to the vP-periphery as in declarative contexts when it bears the relevant features.

Tying these studies together, it is unclear how Bocci & Pozzan's (2014) results could be captured by a purely feature-based account in the sense of Sorace (2000) or Tsimpli et al. (2004). To make this fully

transparent, suppose we follow Belletti (2001, 2004) in saying that postverbal subjects at the vP-periphery in Italian are specified for a discourse (i.e., interpretable) feature like [+ New information focus], whereas preverbal subjects are not. As English postverbal subjects in the relevant contexts are unspecified for such features given the language's non-null-subject status, we might then suggest that under pressure from the L2, postverbal subjects become optionally unspecified for the relevant feature(s) (e.g., [± New information focus]). While this could capture Bocci & Pozzan's (2014) matrix-question results, without further assumption this would wrongly predict attrition in embedded contexts. Without further assumption, this pattern additionally cannot be captured by more recent processing-based accounts such as that suggested Sorace (2011) which assume a special processing status for overt pronouns. Instead, it seems that we must make reference to surface overlap (i.e., conflict) between the word orders of the L1 and L2 in order to account for the infelicitous overextension of discourse-conditioned forms in Bocci & Pozzan (2014). Tying this idea back into Gürel's (2004) ATH-couched account, this could partially solve our observed asymmetry; only in cases of direct conflict between the L1 and L2 do we expect overextension of discourse-conditioned forms such as overt pronouns or post-verbal subjects. This still leaves the source of the increased biases (e.g., null pronouns) unaccounted for. However, given that increased interpretive biases appear to sometimes surface even where we might expect overextension due to conflict (i.e., preverbal subjects in Greek, Tsimpli et al. 2004 and overt pronouns in Spanish, Martín-Villena 2023), we might speculate that increased interpretive biases are the default under attrition, resulting in inconsistent results only where these two general attrition patterns conflict. This is an idea that we will return to and develop in the general discussion.

2.3.3.3 Attrition and production

With the studies above in mind, it is worth noting that the studies so far have predominantly investigated interpretative biases. However, attrition has also been observed to affect the production of null and overt pronouns in null-subject languages. In addition to their picture selection tasks, Martín-Villena (2023) also presented a corpus study to investigate the oral production of null and overt subjects by the same participants (i.e., functional monolinguals, instructed bilinguals, and immersed bilinguals) as in their main interpretive task reported above. For this, they asked participants to watch two silent video clips. After watching each clip as many times as they like, participants were instructed to record themselves narrating the events of that clip in Spanish for someone who had not seen it. The first video clip[25] included only one animate character, Charlie Chaplin. This was done to explore the realisation of topic continuity where the main character is maintained across clauses for an extended period. The second[26] contained six animate characters with different genders. This was included to explore the role of intervening antecedents.

From these recordings, the author annotated and extracted all instances of a 3rd person singular subject in a topic continuity context. These subjects were classified by form (e.g., null pronoun, overt pronoun, proper name, common NP) and syntactic configuration (i.e., intersentential, coordination, or subordination). The author also considered distance from the last 'textual' (i.e., unambiguous full DP) or 'cognitive' (i.e., any DP, including pronouns) mention of the analysed subject as well as the number of proceeding and intervening potential antecedents (regardless of the DP's position within the clause). For that, they considered the four CPs prior to the analysed subject.

Synthesising their results, the author found that although all three groups preferentially produce null subjects in topic continuity contexts, both bilingual groups produce significantly fewer null pronouns than the control group (control: 96.7%; instructed: 95.1%; immersed: 92.9%). However, this trend appears to be driven exclusively by participants' behaviour in the second clip which includes multiple characters. When the first clip (which had a singular character) is analysed in isolation all three groups produce null pronouns at

[25]Link to clip: https://www.youtube.com/watch?v=7xgUdqT6m5A
[26]Link to clip: https://www.youtube.com/watch?v=4QkTNJFhu-g

ceiling (99.2%, 99%, and 98.2% respectively) with no significant global differences between them.

With regard to the distance of the last textual mention of the current topic, the three groups pattern together in mostly producing null subjects within the first two CPs. At a distance of 3 CPs, however, immersed bilinguals were found to produce significantly more overt subjects than the control group. Beyond that point, this difference from the control group was also significant for the instructed bilinguals. When cognitive mentions were considered, the difference between the control and immersed group was found to be significant already at the first CP. For the instructed group a similar difference from the control group was found only when the narration of the second clip was considered in isolation.

As for syntactic contexts, instructed bilinguals were not found to differ from the controls in any context (i.e., coordination, subordination, or intersentential). However, the immersed bilinguals produced significantly fewer null subjects than controls in all three. Additionally, this effect was noticeably larger in intersentential contexts (coordination: 97.1% vs 98.6%; subordination: 94.2% vs 96.6%; intersentential: 81.5% vs 93.2%).

Regarding the role of extra antecedents, when multiple potential antecedents (i.e., characters) preceded the relevant subjects in the discourse contexts, all three groups produced overt subjects at a greater rate. However, again the immersed bilinguals produced significantly more overt subjects than controls. Moreover, this pattern held regardless of whether there was only a single activated antecedent (i.e., there is no ambiguity) or there were 3. For the instructed bilinguals, the increased rate of overt subjects when compared against the control group was only significant in 3 antecedent contexts. Similarly when considering direct linear intervention all three groups produced more overt subjects when another potential referent intervened between the subject and its intended antecedent. However, both bilingual groups were found to produce significantly more overt subjects than the control group when there was an intervening antecedent. For the immersed bilingual, this pattern was significant even when no intervenor was present.

As such, the author presents ample evidence for the overproduction of overt subjects by the experimental groups in a variety of syntactic and discursive contexts. Moreover, for the immersed group, this overextension is attested even when ambiguity resolution is not at stake (i.e., when there was only one active potential antecedent). Integrating the observation that no group-level differences were observed when the narratives for the first video clip were analysed in isolation, however, Martín-Villena (2023) suggest that the overextension of overt pronouns in topic continuity contexts is not a universal feature of attriters. Rather, they suggest that their results relate to the cognitive complexity of retelling the second video clip. In that version of the task, different characters 'took a leading role' at various points. As such, even when no other antecedents were active in the preceding CPs, participants had to consider the larger context of the story. As a result, it is not unreasonable to assume that the retelling of the second video clip is more cognitively demanding than the retelling of the first. Tying back into the IH then, we can say that Martín-Villena's (2023) production task provides two supporting pieces of evidence, the aforementioned overproduction of overt subjects and the restriction of this effect to the more cognitively demanding version of the task.

With regard to the proposed production-comprehension asymmetry, it is worth highlighting that Martín-Villena's (2023) production results contrast with their interpretive and processing results from the same participants. On the one hand, it is clear that native Spanish speakers undergoing attrition overproduce overt pronouns in topic continuity contexts, particularly in the cognitively more demanding version of that task. Moreover, in no context does the author report an increased use of null subjects in topic continuity by the experimental groups. On the other hand, we have argued that Martín-Villena's (2023) interpretive results from the same participants could be used to argue either for a relaxation or strengthening of the non-SpecIP bias for overt pronouns.

2.3.3.4 Attrition and online processing

Moving beyond offline interpretation and production, attrition of pronominal biases in null-subject languages has also been studied using online psycholinguistic measures. As an initial study in this direction, Chamorro et al. (2015a) conducted an eye-tracking-while-reading experiment with disambiguated pronominal items similar to those used in Tsimpli et al. (2004). For that experiment, they employed both online and offline measures. This was done to investigate whether attrition is best accounted for in terms of changes at the level of representations (e.g., Tsimpli et al. 2004) or processing (Sorace 2011). As the authors hypothesised that attrition affects processing but not representations, they expected any attrition effects to be restricted to their online eye-tracking measures.

To test that hypothesis, the authors recruited a group of native (Iberian) Spanish speakers (N = 24) who had recently arrived in the UK (mean length of residency = 7.96 weeks) and used their L1 more frequently than their L2 to serve as a control group. They also recruited an experimental group of native (Iberian) Spanish speakers (N = 24) who had been living in the UK for a minimum of 5 years (mean length of residency = 7 years) and were 'near-native' speakers of their L2 which they reported using more frequently than their L1.

In their experiment, participants were presented biclausal sentences in the order matrix-embedded where the embedded CP contained the relevant pronoun and the matrix CP contained the two possible antecedents. Within these items, they manipulated the pronominal form (i.e., null or overt) and the forced antecedent (i.e., SpecIP or non-SpecIP). To force particular readings, the two possible antecedents in the matrix CP always differed in number (counterbalanced *within* items) such that only one could corefer with the singular pronoun in the embedded CP. An example is given in (33) with vertical bars to indicate the critical region for their eye-tracking measures. During the experiment items appeared alone on the screen. After participants had read each item, they were asked to judge how natural the sentence sounded on a scale from 1 ('not natural at all') to 5 ('totally natural').

(33) a. *La madre saludó a las chicas cuando | (ø / ella) cruzaba | una calle con*
the mother.SG greeted DOM the girls.PL when | (*pro* / she) crossed.SG | a street with
mucho tráfico. [Spanish]
much traffic.

'The mother greeted the girls when she crossed the street with a lot of traffic.'

b. *Las madres saludaron a la chica cuando | (ø / ella) cruzaba | una calle con*
the mother.PL greeted DOM the girl.SG when | (*pro* / she) crossed.SG | a street with
mucho tráfico.
much traffic.

'The mothers greeted the girl when she crossed the street with a lot of traffic.'

Given the PAH, one would expect faster reading times and higher acceptability judgements from native (Iberian) Spanish speakers when null pronouns are coreferential with the possible antecedent in SpecIP. For overt pronouns, the opposite pattern would be expected.

Starting with the offline acceptability judgement task, results for their control group indicated an interaction of *pronoun* and *antecedent* which was significant by subjects (ANOVA). This surfaced as significantly higher acceptability for non-SpecIP disambiguated overt pronouns than SpecIP disambiguated ones (3.60 and 3.26 out of 5 respectively). No effect was observed within the null pronouns (non-SpecIP = 3.61; SpecIP = 3.72). For the experimental group, the interaction was significant by subjects and items. This was again driven by higher acceptability for overt pronouns when disambiguated toward the non-SpecIP antecedent

(non-SpecIP = 3.45; SpecIP = 2.91) with no effect difference within the null items (non-SpecIP = 3.54; SpecIP = 3.64). When the groups were compared no three-way interaction with *group* was observed. As such, the authors interpret these results to indicate two findings: (i) that their offline acceptability judgement task is tapping antecedent preferences stemming from the PAH at least for overt pronouns and (ii) that this offline sensitivity to the PAH is not affected by attrition.

Turning to their online eye-tracking measures, the authors considered three common measures. The first pass (FP) measures the sum of all fixations from the first time a region is entered from the left until the first time it is exited either to the left or the right. The go past (GP) measure is the duration of all fixation from the first time a region is entered from the left until it is exited on the right (i.e., a progression). That measure may include fixations to the left of the current region (i.e., regressions). The total time (TT) measure is the sum of all fixations in a region over the course of an item.

For their control group, they reported a *pronoun* by *antecedent* interaction at the critical region that was significant by items and subjects in the FP measure, and by subjects in the GP and TT measures. Similar to their offline results, *post hoc* tests indicated that this was driven by items with overt pronouns which were read more quickly when disambiguated toward non-SpecIP antecedents. For their experimental group, the *pronoun* by *antecedent* interaction was not significant in any measure. Moreover, when the groups were compared they observed a significant three-way interaction with *group*. From this, they concluded that speakers of (Iberian) Spanish are generally sensitive to PAH biases in the online processing of overt pronouns, but that speakers undergoing attrition are not. Integrating their online and offline results then, they take this as evidence to support their hypothesis that attrition affects the online processing of the relevant structures, not their underlying representations.

Tying this into the larger literature, Chamorro et al.'s (2015a) eye-tracking results indicate that the offline overextension of overt pronouns attested in Tsimpli et al. (2004), Gürel (2004), and Gürel & Yılmaz (2011) extends to reduced sensitivity to biases in online processing. However, no increased matrix subject bias was observed for null pronouns despite this pattern being attested in the offline interpretation biases of other null-subject languages (Tsimpli et al. 2004, Gürel 2004, Gürel & Yılmaz 2011). Curiously, however, Chamorro et al. (2015a) found no evidence for a clear SpecIP bias with null pronouns in any measure, online or offline. Moreover, when Chamorro (2018) constructed ambiguous versions of the items in Chamorro et al. (2015a) by removing the number disambiguation, native Spanish speakers (N = 24) were found to exhibit a significant non-SpecIP interpretive bias for overt pronouns but no clear bias for items with null pronouns. This is despite other studies with Iberian Spanish speakers finding a clear SpecIP bias for null pronouns in different items (e.g., Martín-Villena 2023). As such, the lack of effects may relate to the particular items.

Around the same time, Kaltsa et al. (2015) presented a self-paced listening task to explore the online processing of pronominal biases in Greek. For their experiment, the authors recruited an experimental group of L1-Greek speakers who had grown up in Greece but then immigrated to Sweden in adulthood (N = 25) and acquired Swedish as an L2. Relevant to the present discussion, Swedish like English is a non-null-subject language (Platzack 1987).[27] On average those participants had lived in Sweden for 31 years (min = 25 years) and had a mean age of 59 years (range: 55 - 65). They also recruited an age-matched control group of monolingual Greek speakers (N = 18).[28]

Critical items (which appear to have been adapted from the Greek stimuli in Tsimpli et al. 2004) always consisted of an initial matrix CP which contained two possible antecedents followed by an embedded CP which contained a globally ambiguous pronoun. In half of the items, the pronoun was null, and in half it was overt. Sentences were presented in a self-paced listening task (with boundaries indicated by vertical bars in

[27] Kaltsa et al. (2015) set aside the issue of demonstrative pronouns in Swedish which, like overt pronouns in Greek, exhibit a non-subject bias in intrasentential contexts. This is because the use of such demonstrative pronouns is marked and restricted to written or high registers.

[28] That study also two additional groups: a second experimental group consisting of HSs of Greek, and a second age-matched control group. For the comparison of HSs and potential attriters please see the original paper.

34 and 35). As soon as a participant started listing to the first segment, they were presented a single image which corresponded to one of the three possible interpretations of the ambiguous pronoun. An example of the three possible images associated with (34) is presented in Figure 2.6. After listening to the final segment, participants were presented a question mark in the place of the image. They then had to indicate if the image they had seen matched the sentence they had heard.

(34) *I γiaγia | xeretise | tin kopela | otan | afti | pernuse | to δromo.* [Greek]
the old.lady | greeted | the girl | when | she | crossed | the street

'The old lady greeted the girl when she crossed the street.'

(35) *O papus | milise | δinata | ston egono tu | otan | ø δiavaze | ena vivlio.*
the old.man | spoke | loudly | to | grandson his | when | *pro* read | a book

'The old man spoke loudly to his grandchild when he read a book.'

(a) subject reading (b) object reading (c) other reading

Figure 2.6: Example images used in Kaltsa et al. (2015) for the item in (34).

From this experiment, the authors present three different measures: (i) acceptance rates for the picture judgement portion of the study, (ii) decision times for the judgements, and (iii) listening times for windows within the embedded CP. We do not consider the third of these as it is not clear what the original authors predicted, nor is it clear how listening times could relate to the event depicted in the images in the items with overt pronouns. This is because, in those items, listening times were reported for the window containing only the overt pronoun (window 5 in 34). As that window appeared separately from the embedded verb, items were still temporarily compatible with other embedded predicates such as 'waits on the pavement'.

For interpretations of overt pronouns, the authors predicted the control group would exhibit a non-subject bias which would be weakened under attrition resulting in a greater acceptance of subject images in the experimental group. For null pronouns, both groups were predicted to prefer subject interpretations with no difference between them. With regard to decision times, the authors predicted no group-level difference for null pronouns given that the interpretive bias for null pronouns was predicted to be unaffected by attrition. For overt pronouns conversely, the experimental group was expected to respond more slowly overall given that group was predicted to exhibit weaker interpretive biases.

Abstracting away from some of the less relevant aspects of their results, analysis of the picture judgements for items with overt pronouns indicated that both groups accepted object images significantly more frequently than subject or object images and the experimental group additionally accepted subject images significantly more frequently than 'other' images (experimental: subject: 58%, object: 80%, other: 32%; control: subject: 31%, object: 66%, other: 25%). Between-group pairwise comparison also indicated the experimental group accepted subject interpretations significantly more frequently than the control group (with no difference for object or other). The authors interpret this to indicate a relaxation of the non-subject bias for overt pronouns.

For decision times, the control group gave responses to object images significantly faster than subject ones which in turn were responded to more quickly than other images. For the experimental group, responses to object images were significantly faster than to subject images or for other images, with no difference between subject and other. Under pairwise comparison, the control group was found to respond significantly more quickly to object and other images with no difference for subject images.

Tying these measures together, results seem to match the authors' predictions. Offline measures indicate a clear attrition effect with experimental participants being more accepting of subject readings despite a continued object bias. Moreover, online measures indicated that the control group was faster than the experimental group in their object and other responses as expected which might suggest that the experimental group is less sure about their interpretive biases. However, we might also note that both groups responded more quickly to object images than subject ones. That is to say, at least for the linguistically introduced referents, both groups' response times were faster when the image matched their preferred interpretation.

For items with null pronouns, the experimental group accepted subject images significantly more frequently than the alternatives (subject: 85%, object: 76%, other: 48%). For the same items, the control did not differ in their ratings for subject and object images although they accepted both object and subject images higher than 'other' images (subject: 77%, object: 66%, other: 33%). However, pairwise comparison of the two groups indicated that they did not differ in their acceptance rates of any image type for null pronouns. As such, given the similar numeric trends, the authors interpreted this to indicate that the SpecIP bias for null pronouns is stable under attrition.

For decision times, pairwise comparison indicated that the experimental group responded more slowly to subject images than object images, which in turn were responded to more slowly than other images. For the control group instead, responses to other images were significantly slower than subject and object images which did not differ from each other. When the two groups were compared pairwise, the experimental group was found to respond significantly more slowly than the control group for both subject and object images, but significantly faster for other images. Thus, integrating the decision times from the null and overt items, it appears that increased decision times under attrition (null and overt) do not appear to relate to the increased flexibility of pronominal interpretations (overt only) *contra* the predictions of Kaltsa et al. (2015). Instead, we might interpret their results to indicate that the attriters were less certain/deterministic in their judgements regardless of whether their end judgements differed from the controls'. Tying this study into the larger literature then, we cannot take the results in Kaltsa et al. (2015) to support Chamorro et al.'s (2015a) conclusion that attrition results in a reduced sensitivity to the PAH in online processing specifically for overt pronouns.

More recently, Martín-Villena (2023) conducted a self-paced reading task with the same participants from their picture selection and production tasks (i.e., functional monolinguals, instructed bilinguals, and immersed bilinguals). For that task, stimuli were adapted from those in the picture selection task.[29] To re-iterate, stimuli always consisted of a matrix CP that contained the possible antecedents followed by an embedded CP with the ambiguous pronoun. To keep the number of windows (i.e., 7) identical between null and overt items, an adverb was added to the matrix CP of the items with null pronouns. These items were then presented phrase-by-phrase (with window boundaries indicated by vertical bars in 36) in a non-cumulative fashion. Before each item, participants were presented an image alone on their screen. This corresponded to either a SpecIP or non-SpecIP reading of the pronoun in the sentence that they were about to read. As an example, the relevant images for the sentence in (36a) are provided in Figure 2.7. After they had examined the image, participants pressed the space bar to see the first window. Regardless of the window the text always appeared immediately below the image which remained on the screen. After they had read the entire item, a comprehension question then appeared below the image asking whether the image on the screen matched

[29] For discussion regarding changes to *cuando* ('when') and *mientras* ('while') items, please see the original thesis.

the sentence that they had read (yes or no).

(36) a. *El abuelo* | *habló* | *rápido* | *al nieto* | *mientras* | *ø leía* | *el libro.* [Spanish]
the grandad talked fast to.the grandson while *pro* read the book

'The grandfather spoke quickly to the grandson while he was reading the book.'

b. *La madre* | *besó* | *a la hija* | *mientras* | *ella* | *se ponía* | *el abrigo.*
the mother kissed DOM the daughter while she REFL put.on the coat

'The mother kissed the daughter while she was putting on the coat.'

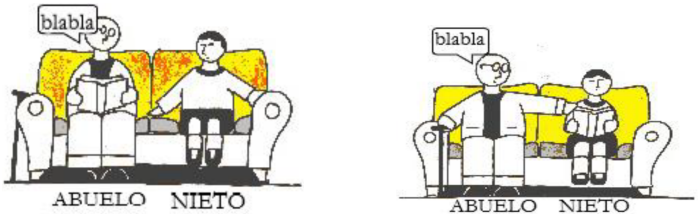

(a) SpecIP reading (b) Non-SpecIP reading

Figure 2.7: Example of the images used for the self-paced reading task in Martín-Villena (2023) for the item in (36a).

In their analysis of this task the author considers three measures: (i) the combined reading time for all windows after the adverbial (e.g., *mientras* – 'while' or *cuando* – 'when'), (ii) the response time to the comprehension question, and (iii) the responses themselves. Additionally, as items with overt pronouns trivially included an extra window corresponding to the pronoun, they analysed reading times for null and overt items separately. For this study, the author predicted that participants would read the embedded CP more slowly when the event depicted in the co-occurring image conflicted with their preferred pronominal interpretation. Given that the control group was expected to exhibit biases in line with the PAH, one would expect reading times for items with null pronouns to be faster when presented with a SpecIP image than a non-SpecIP one, and the inverse for items with overt pronouns. If one predicts a reduced sensitivity to the non-SpecIP bias for overt pronouns under attrition, the two bilingual groups would be expected to differ from the control group in items with overt pronouns only.

Starting with reading times for null pronominal items, the effects of *group* or *bias* (i.e., subject or object image) were non-significant, as was their interaction. Additional modelling in which *BLP score* and *working memory score* were entered as simple predictors or potential interaction terms with *bias* did not improve model fit (although see the earlier discussion on the limitations of these measures in a model with *group*). For reading times of items with an overt pronoun, the effect of *bias* was found to be significant for the control group, but not in either experimental group. However, *contra* Martín-Villena (2023) we should refrain from interpreting this as indicating a difference in sensitivity to the PAH given that the author did not comment on the significance of the *group* by *bias* interaction, nor did they present *post hoc* comparisons between the groups. Similar to the model for reading times of the items with null pronouns, the inclusion of *working memory score* and *BLP score* did not improve model fit.

Turning to decision times, modelling considered overt and null pronominal items together. For that model, the author only reported a significant three-way interaction of *group*, *bias*, and *pronoun*. Following up on this, *post hoc* tests indicated that for overt items, all three groups responded more quickly in the object bias condition than the subject bias one. That is to say, all participants were faster to respond when the image matched the predicted PAH bias. Conversely, for items with null pronouns, only the two bilingual groups exhibited a significant effect of *bias* responding more quickly in the subject bias condition. For the control

group, this effect was non-significant but trended in the same direction. Inclusion of *working memory score* and *BLP score* as potential interaction terms did not improve the model.

For the analysis of responses to the comprehension questions, responses were first coded as ± PAH-compliant. For example, if a participant saw an item with a null pronoun (36a) and a subject-biasing image (Figure 2.7a), a 'yes' response was coded as 1 and a 'no' response was coded as 0. If instead, the participant saw an object-biasing image for the same sentence (Figure 2.7b), a 'yes' response was coded as 0 and a 'no' response was coded as 1. The ± PAH-compliant responses were then subjected to a mixed-effect logistic regression. That model included effects of *pronoun* and *bias* as well as their interaction. *Group* was not included in the final model, nor was *BLP score* or *working memory score*. This was because no higher-order interaction was found to significantly improve the model. As such, this model cannot provide evidence for attrition.

Nonetheless, the author presented further models in which they considered whether *LoR/LoI* affect the *pronoun* by *bias* interaction when the bilingual groups are considered in isolation. For the instructed bilingual group, this did not improve model fit and no values are reported. For the immersed bilingual group however, the author reports a significant three-way interaction of *bias*, *pronoun*, and *LoR*. Although *post hocs* are not presented, the author reports that the effect of *LoR* was significant in all the bias-pronoun combinations save for null pronouns in combination with an object image. They take this as evidence in support of the IH for two reasons. First, increased residence in the L2 is associated with more PAH-non-compliant responses for overt pronouns in combination with a subject image. Second, the same does not hold for null pronouns in combination with an object image. However, by their own account, *LoR* also affected null pronouns in combination with subject images, and overt pronouns in combination with object images. In the case of the former, their Figure 60 (page 275) clearly indicates a positive effect. That is to say, participants gave more PAH-compliant responses as their length of residence increased. This is incompatible with the IH, but convergent with the increased L1 bias pattern observed in many of the above studies. As such, we should be cautious of interpreting these results without the unreported *post hocs*, but should one choose to interpret these results, it should be recognised that they are mixed.

In summary of the above subsection, there is still relatively little work on attrition and the PAH in online processing. Moreover, of the available studies, only one (i.e., Chamorro et al. 2015a) provides unambiguous evidence for a relaxation of the non-SpecIP bias for overt pronouns under attrition while observing no increased bias for null pronouns. Although in that particular case, we have suggested that the lack of an effect may relate to the stimuli employed.

2.3.3.5 Attrition and re-exposure effects

Another issue that we must take into consideration when discussing attrition, particularly from an online perspective, is re-exposure. This is because some researchers have argued that re-exposure to the L1 may lead to a mitigation or reversal of attrition effects.

Circling back to Chamorro et al. (2015a), those authors provide some initial evidence to this effect. Recall that they conducted an eye-tracking-while-reading experiment using items like in (33, partial) in which they manipulated the pronominal form (i.e., null or overt) and the forced antecedent (i.e., SpecIP or non-SpecIP). In addition to their controls and potential attriters, the authors recruited a third group of native (Iberian) Spanish participants (N = 24) for that experiment. Those 're-exposees' were similar to the potential attriters in that they had been living in the UK for a minimum of 5 years (mean = 5.83 years) and had achieved 'near-native' proficiency in their L2 which they spoke significantly more frequently than their L1. However, unlike the potential attriters, the re-exposees were tested shortly after returning from their winter holidays in Spain (minimum duration = 7 days, mean = 13.08).

(33, partial) a. *La madre saludó a las chicas cuando* | *(ø / ella) cruzaba* | *una calle*
the mother.SG greeted DOM the girls.PL when | (*pro* / she) crossed.SG | a street
con mucho tráfico. [Spanish]
with much traffic.

'The mother greeted the girls when she crossed the street with a lot of traffic.'

Offline results from the re-exposees (i.e., acceptability judgements after each item) indicated an interaction of *pronoun* and *antecedent* which was significant by subjects. This was due to lower acceptability ratings for overt pronouns when they were forced to corefer with the potential antecedent in matrix SpecIP. No effect was observed within the null pronouns. This was in line with the results from both the controls and potential attriters. Indeed, the re-exposees were not found to differ significantly from either group.

As for eye-tracking measures in the critical region (indicated by the vertical bars in 33, partial), the interaction of *pronoun* and *antecedent* was significant in the TT and marginal in the FP measures. This was driven by significantly slower reading times for items with overt pronouns that were disambiguated toward the potential antecedent in the matrix SpecIP. No effect of *antecedent* was observed for items with null pronouns. As such, results from this group pattern more similarly with the results from the control group than the results from the potentially attriting group. The similarity with the control group was supported by the lack of a significant three-way interaction with *group* in any measure when the two groups were compared. From this, the original authors conclude that attrition is sensitive to changes in the input such that even relatively short periods of re-exposure to the L1 may lead to mitigation. However, when the re-exposees and potential attriters were compared, three-way interaction with *group* was again not significant in any measure. Thus, although their interpretation is indeed a possible one, we should also keep in mind that it rests on null results and therefore should not be taken as conclusive.

Nonetheless, convergent results have since been reported for a case study of a single L1-Bulgarian (consistent null-subject language), late L2-German (partial null-subject language) speaker (Köpke & Genevska-Hanke 2018). In spontaneous speech recorded after 12 years of living in Germany, the authors observed that the case study participant produced overt pronominal subjects in their L1-Bulgarian at a significantly higher global rate (47% of pronominal subjects) than 10 L1-Bulgarian speakers still living in their home community (average 27%). As such, the case-study participant appears to be undergoing attrition at the initial record. Tying in with the discussion of re-exposure however, they also recorded the same speakers after two weeks of holiday in their L1 community. While in their L1 community, the case study participant was found to produce overt pronominal subjects at a rate similar to the speakers who remained in their home community (34%), with no statistically significant difference between them. Moreover, when the two recordings of the case-study participant were compared against each other, the two recordings were found to differ significantly.[30] The authors take this to indicate that re-exposure to the L1 had attenuated the attrition effect reported in the Germany recording, thus providing support for Chamorro et al.'s (2015a) conclusion.

Additionally, those authors argue that re-exposure need not require travel to the L1. This is because the same case study participant was recorded 5 years later. As before, recordings were taken both in Germany and in Bulgaria again after 2 weeks of holiday. In that second set of recordings, their production of overt pronouns both in Germany (29%) and in Bulgaria (24%) was in line with the Bulgarians living in their home community (i.e., attrition was no longer attested). The authors speculate that this may relate to the fact that between the first and second sets of recordings, the speaker married another L1-Bulgarian speaker. As a result, there was a noticeable change in their pattern of language use. At the initial testing point, the speaker had 'extremely limited contact' with her L1 and was deemed to be dominant in her L2-German. At the second time point, however, the speaker was reported to have a 'more balanced pattern of use for the two languages.'

[30] Although Köpke & Genevska-Hanke (2018) provide *p*-values for the comparisons of the case-study participant against the controls, they do not report their stats for the comparison of the relevant two recordings of the case-study participant.

To the best of our knowledge, the only other non-case study to explore re-exposure effects on attrition on syntactic biases is reported by Gargiulo & van de Weijer (2020). They conducted a sentence interpretation task to tap pronominal biases in a group of potential attriters before and after re-exposure to their L1 Italian. For this, they recruited two groups of participants. Their control group (N = 21) consisted of native Italian speakers still residing in Italy whereas the experimental group (N = 20) consisted of L1-Italian L2-Swedish (a non-null-subject language) speakers who had grown up in Italy but had been living in Sweden for a minimum of 7 years prior to testing (mean = 11.83). That group rated their proficiency in Swedish as relatively high, 4.75/5 (where 1 indicated 'very low' and 5 indicated 'very good').

In their task, participants were presented ambiguous items as in (37). These items always appeared alone on the screen in the order matrix-embedded, with the pronoun in the embedded CP introduced by *quando* ('when') or *da quando* ('since'). After reading each sentence, participants pressed the space bar and were presented a comprehension question alone on the screen. This always asked about the referent of the ambiguous pronoun with the two possible responses from the preceding linguistic context (38).

(37) *Monica ha discusso molto con Antonella da quando (ø / lei) è tornata da Parigi.*
Monica has discussed much with Antonella from when *pro* / she is returned from Paris

'Monica has discussed a lot with Antonella since she came back from Paris.' [Italian]

(38) *Chi è tornata da Parigi?* [Italian]
who is returned from Paris

'Who came back from Paris?'

a. Monica

b. Antonella

To explore re-exposure effects, the experimental group participated in the same task twice: once prior to their summer vacation to Italy (min 11 days, mean = 23.2 days), and once as soon as possible after their return to Sweden (mean = 2.95 days). For consistency, the control group also participated twice (min time between sessions 20 days, mean = 22 days). This was done to account for any effects of task repetition.

Results indicated that both groups, at both time points, exhibited PAH-compliant biases of both pronominal types: for items with null pronouns, participants preferentially selected SpecIP responses, for items with overt pronouns participants preferentially selected non-SpecIP responses (Table 2.1). When analysing the data, the authors chose to only report a simplified model with two-way interaction terms as an initial model with the three-way interaction of *session*, *pronoun*, and *group* indicated that the three-way interaction was not significant. From their simplified model, the authors report a significant two-way interaction of *pronoun* and *session*. Based on visual inspection in the absence of *post hoc* tests, the authors interpret this as an increased SpecIP bias for both groups in the second session. The authors also report a significant interaction between *pronoun* and *group* (i.e., an attrition effect). This was interpreted as a weaker SpecIP bias for null

	Control		Experimental	
Session	Null	Overt	Null	Overt
First	78%	9%	74%	16%
Second	88%	11%	80%	13%

Table 2.1: Percentage of SpecIP responses by group and pronoun in both sessions adapted from Gargiulo & van de Weijer's (2020) Table 1.

pronouns in the experimental group (despite a similar-sized trend in the overt pronouns). However, this was again based on visual inspection of the data in the absence of *post hoc* tests. Thus, they arrive at two relevant conclusions. First, despite the observed difference between the pre- and post-vacation testing sessions, this is not attributable to L1 re-exposure. Rather this is taken to be a task-repetition effect given similar results were observed in the control group. Second, the authors conclude that attrition in their sample affected null pronouns resulting in a weaker SpecIP bias *contra* Tsimpli et al. (2004) and the emergent pattern that we have argued for above.

2.3.3.6 Summary of previous results

Summarising our review of the previous literature on attrition in null-subject languages, we have argued for two general patterns. The first of these is the often discussed relaxation of the discourse restrictions (particularly for overt pronouns) on syntactic options offered by null-subject languages, with unambiguous evidence from interpretation (Tsimpli et al. 2004, Gürel 2004, Kaltsa et al. 2015), production (Köpke & Genevska-Hanke 2018, Martín-Villena 2023), and online processing (Chamorro et al. 2015a). Moreover, this effect appears to be modulated by re-exposure to the L1 (Chamorro et al. 2015a, Köpke & Genevska-Hanke 2018), and is sensitive to (surface) overlap between the L1 and L2 (Gürel 2004, Bocci & Pozzan 2014). The second pattern we have argued for is an apparent increased sensitivity to interpretive biases already present in the L1 which cannot be directly attributed to the influence of the L2. This was found to affect the interpretation of null pronouns (Tsimpli et al. 2004, Gürel 2004, Gürel & Yılmaz 2011), overt pronouns (Martín-Villena 2023), and the interpretation of preverbal non-pronominal subjects (Tsimpli et al. 2004). This second pattern raises the question of whether a similar increased sensitivity to L1 interpretive biases may be observed in ambiguities which are (i) unrelated to the null-subject status of these languages and (ii) not argued to be conditioned by interpretable discourse features. To explore this idea, the next section introduces another phenomenon, (pseudo-)relative clause attachment ambiguities.

2.4 (Pseudo-)relative clause parsing ambiguities

Upon encountering an incoming string, the hearer/reader must decode the physical signal and arrive at a syntactic representation, a process known as syntactic parsing. Sometimes, the signal may be compatible with multiple representations. In such cases, many researchers have assumed that the parser relies on principles of computational efficiency such as locality or minimal structure to resolve ambiguity and reduce processing cost (e.g., Kimball 1973, Frazier 1978, De Vincenzi 1991). This has led to non-trivial interest in the parsing of RC attachment. To see why, consider (39). In that example, the bracketed RC may modify to either NP_1 or NP_2. These two readings are known as high attachment (HA) and low attachment (LA) readings respectively due to the relative height of the RC as sketched in (40).

(39) a. Someone shot the $maid_1$ of the $actress_2$ [that was standing on the balcony].

b. *Alguien disparó contra la criada$_1$ de la actriz$_2$ [que estaba en el balcón].* [Spanish]
someone shot against the maid of the actress that was at the balcony

'Someone shot the maid of the actress that was standing on the balcony.'

Cuetos & Mitchell (1988)

This presents an example of an attachment ambiguity. Should there be a universal bias that favours local attachment (e.g., LATE CLOSURE), we would expect speakers of any language with RCs to prefer LA readings in such contexts. Since Cuetos & Mitchell (1988) however, it has been well known that speakers of different languages exhibit contrasting biases in their interpretation of sentences like (39). Whereas speakers of English

(40) a. High Attachment

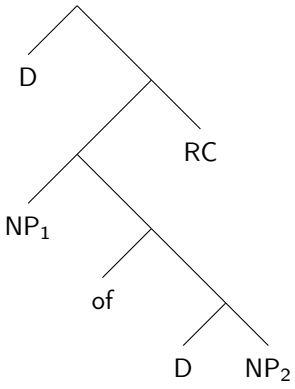

b. Low Attachment

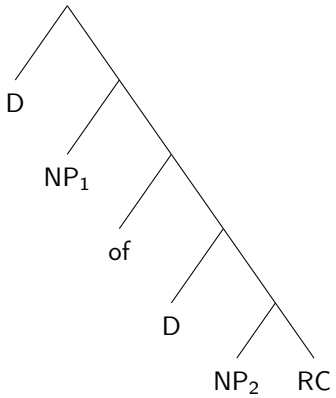

have been shown to exhibit the expected LA bias, speakers of Spanish have been repeatedly shown to exhibit a HA bias (i.a., Cuetos & Mitchell 1988, Carreiras & Clifton 1993, 1999, Fernández 2003, Dussias 2003). This has led to some researchers rejecting locality as a universal parser bias (e.g., Mitchell & Cuetos 1991: et seq.) as well as a classification of languages as either LA languages (e.g., English: Cuetos & Mitchell 1988, Fernández 2003, i.a.; Chinese Shen 2006) or HA languages (e.g., Spanish; French: Zagar, Pynte & Rativeau 1997; Greek: Papadopoulou & Clahsen 2003).

More recently, Grillo (2012) and Grillo & Costa (2014) have suggested that this crosslinguistic difference in parser biases may only be apparent. Rather, they point out that there is a hidden structural difference which must be taken into account. Namely, many of the languages that have been previously reported to exhibit a HA bias also admit PRs, whereas languages that exhibit a LA bias do not. PRs, despite being string identical to true RCs, display a number of semantic and syntactic differences. Those authors also observed that in the contexts that have been previously studied in the psycholinguistic literature, PRs are only compatible with HA. Thus, they suggest that once this syntactic confound is accounted for we can maintain locality as a universal parsing bias.

Before delving into the specifics of the PR account[31] and its predictions regarding offline interpretive biases and online processing (section 2.4.1) or how these biases are affected by attrition (section 2.4.3), it is worth pausing to highlight why attachment ambiguities are relevant to the discussion above. First, attachment ambiguities offer a clear example of interpretive biases unrelated to the licensing of null subjects via agreement on T. Second, the interpretive biases that we will focus on are not argued to be conditioned by discourse features. Rather, they are assumed to derive from computational efficiency at the level of the parser.[32] Consequently, by exploring how attrition affects attachment ambiguities, we can isolate the role of interpretive/processing biases from the syntax-discourse interface. That in turn can help us better understand what makes certain structures more vulnerable to attrite. If there is something about the more interface structures that causes them to be more susceptible to attrite (e.g., in the case of null and over pronouns, that the integration of discourse information is more cognitively costly than the processing of more language internal information, Sorace 2011), then attrition should affect pronominal reference and attachment

[31] Within this thesis, we will focus on the PR account to the exclusion of alternatives like the Tuning Hypothesis (Mitchell & Cuetos 1991) or Construal Theory (Frazier & Clifton 1995). This is because there is general agreement that earlier accounts cannot fully capture the observed crosslinguistic variation. Additionally, the PR account allows us to generate useful language internal predictions. These can, in turn, be used to make novel predictions about attrition which we will exploit in our experiments.

[32] This is not to suggest that the processing of such structures does not require information from any other domain. The processing of the relevant strings will trivially require things like lexical access and working memory. Rather, our point is that along a spectrum of more external and more internal interface phenomena, we take intrasentential pronominal resolution as an example of a more external interface phenomena and 'RC' attachment ambiguities as an example of a more internal interface phenomena even if it still requires the integration of other non-syntactic information.

ambiguities differently. However, if attrition leads to (i) a relaxation of certain biases under surface-level overlap with the L2 as well as (ii) a general increase in interpretive biases already present in the L1 as argued above, then we would expect similar effects in the parsing of attachment ambiguities despite the difference in interface status.

2.4.1 The PR account

Although string identical to true RCs, PRs exhibit a number of syntactic and semantic differences (i.a., Radford 1975, Cinque 1992, Casalicchio 2013). On the semantic side, RCs denote properties of the NPs they modify, but PRs denote situations in events (Moulton & Grillo 2015). Therefore, the RC reading in (41a) tells us something about the relevant boy (i.e., that he ate ice cream) whereas the PR reading in (41b) tells us what Gianni saw (i.e., an event of eating in which a particular boy was the agent).

(41) a. *Gianni ha visto* $[_{DP}$ *il* $[_{NP}$ *bambino* $[_{CP}$ *che mangiava il gelato]]].* [Italian]
Gianni has seen the boy that ate the ice.cream
'Gianni saw the boy that was eating the ice cream.'

b. *Gianni ha visto* $[_{PR}$ $[_{DP}$ *il* $[_{NP}$ *bambino]]* $[_{CP}$ *che mangiava il gelato]].* [Italian]
'Gianni saw the boy eating ice cream.'

On the syntactic side, when a PR occupies the complement of a predicate of perception as in (41b), the embedded CP is a sister to the DP. This means that when a string with a complex DP like the ones in (39, 42) is parsed as containing a PR, only the structurally higher of the DPs may c-command into the embedded CP and be interpreted as its subject. That is to say, PR parses force HA in this context.[33] Conversely, RCs are adjuncts within the DP. As such, they may freely attach to either NP_1 or NP_2 in strings like (39, 42).

(42) *Gianni ha visto il figlio del medico che mangiava il gelato.* [Italian]
Gianni has seen the son of.the doctor that ate the ice.cream
'Gianni saw the son of the doctor (that was) eating the ice cream.'

With the above in mind, it is not enough for the parser to simply decide where to attach the embedded 'RC' when presented with such strings. Rather, it must additionally decide whether the embedded CP should be parsed as a RC or a PR. To guide the resolution of this ambiguity, Grillo & Costa (2014) propose the PR-First Hypothesis (PRFH) in (43).

(43) PRFH: When PRs are available, everything else being equal, they will be preferred over RCs.
(Grillo & Costa 2014)

This principle then interacts with locality to derive the crosslinguistic pattern. In languages like Italian or Spanish, the PRFH straightforwardly accounts for the observed HA bias. Whenever a string is ambiguous between a RC and PR parse, the parser will prefer a PR parse, resulting in HA. In languages like English however, the PRFH can only apply vacuously as PR parses are not licit.[34] In that case, the embedded CP must be parsed as a RC, resulting in a true attachment ambiguity. Principles of locality would then obtain, driving the observed LA bias. The interaction of these two principles is sketched in Figure 2.8 with preferred parses noted in boldface.[35]

[33] Although there is no true high/low-attachment ambiguity in such cases, we will continue to call this a 'HA' reading following Grillo & Costa (2014) because the subject of the embedded CP is interpreted as the structurally higher of the relevant NPs.
[34] For the scope of the present thesis, we set aside embedded CPs without an overt complementiser. See Grillo, Costa, Fernandes & Santi (2015) regarding the PR-FIRST HYPOTHESIS and reduced relatives/small clause ambiguities in English.
[35] Although sketched serially, the PRFH is compatible with ranked parallel models.

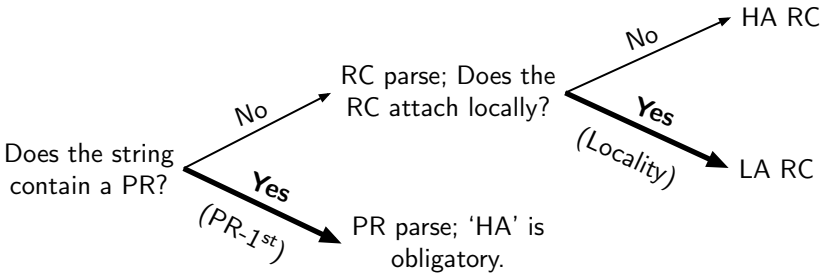

Figure 2.8: Summary of (relevant) universal PR/RC parse options.

2.4.2 Previous (monolingual) work on PRs

As the PR account derives the crosslinguistic pattern from this interaction of principles, we can also use it to make testable predictions language internally. Namely, when PRs are locally blocked in an otherwise 'HA' language, we would expect a LA bias to emerge with *bona fide* RCs. To explore this prediction, Grillo & Costa (2014) conducted a sentence interpretation task in which they manipulated PR availability via the matrix predicates. Although PRs may occupy the complement of a verb of perception such as *vedere* ('see'), they may not follow non-perceptual predicates such as *sposare* ('marry'). This is shown in (44) as the proper name *Maria* blocks a (restrictive) RC reading of the embedded CP, forcing it to instead be interpreted as a PR where grammatical. Changing the *Maria* to a common noun in (45), we can see that RCs are not similarly restricted by the selectional properties of the matrix predicate.

(44) a. *Gianni ha visto Maria che correva.* [Italian]
Gianni has seen Maria that ran

'Gianni saw Maria running.'

b. * *Gianni ha sposato Maria che correva.*
Gianni has married Maria that ran

(45) *Gianni ha sposato la donna che correva.* [Italian]
Gianni has married the woman that ran

'Gianni married the woman that was running.'

Exploiting this restriction on PRs, Grillo & Costa (2014) constructed sentences like (46). Half of the time, items contained a perceptual matrix predicate and half the time they contained a non-perceptual one. They then presented these to native Italian speakers (N = 30) and asked them *who*-questions targeting the subject of the embedded CP. If HA is not a general bias in Italian but is driven by a bias for PRs in compatible contexts as suggested by the PRFH, one would expect a HA bias only in the perceptual items with LA bias in the non-perceptual ones. Consistent with this prediction the authors observed a HA bias (78.6% HA) when items contained a matrix predicate of perception. However when the same items contained non-perceptual predicates, participants preferentially selected LA (24.2% HA), resulting in a significant effect of *predicate type* in their mixed-effects regression. Convergent results have since been reported in other HA languages (e.g., French: Pozniak, Hemforth, Haendler, Santi & Grillo 2019, Spanish: Aguilar & Grillo 2021).

(46) a. *Gianni ha visto il figlio del medico che correva la maratona.* [Italian]
Gianni has seen the son of.the doctor that ran the marathon

'Gianni saw the son of the doctor (that was) running the marathon.'

b. *Gianni vive con il figlio del medico che correva la maratona.*
Gianni lives with the son of.the doctor that ran the marathon

'Gianni lives with the son of the doctor that was running the marathon.'

Moreover, results from English, a LA language, suggest that the findings from HA languages cannot be entirely due to the semantics of the matrix predicate, but rather should be interpreted as an effect of PR (un-)availability. Grillo et al. (2015) conducted a similar experiment with native English speakers (N = 30) which contained translation equivalents of the items in Grillo & Costa (2014) as in (47). If the effect of *predicate type* observed by Grillo & Costa (2014) in Italian were driven by the semantics of the predicates involved, a similar effect would be expected in English. Conversely, if PR availability drove the HA bias in Italian, no such effect would be predicted to surface in English. Grillo et al. (2015) report a small, but significant, effect of semantic facilitation (non-perceptual: 19.5% HA; perceptual: 37.2%[36]). Despite this significant effect, they take their results to be consistent with the PRFH. That is because the English speakers did not exhibit a wholesale shift from LA to HA between non-perceptual and perceptual predicates as was observed in Italian (non-perceptual: 24.2%; perceptual: 78.6%). To support this idea, the authors conducted an additional model[37] in which they directly compared their English results with the Italian results from Grillo & Costa (2014). That revealed a significant interaction of *language* and *predicate* with *post hoc* tests indicating that the effect of *language* was only significant within the perceptual items. Thus, while the English results indicate that the semantics of the matrix predicates play a role, this cannot fully account for the Italian results.

(47) a. Kelly heard the grandma of the girl that was screaming

b. Kelly works with the grandma of the girl that was screaming.

In addition to offline interpretive biases, we can also derive testable predictions regarding online parsing from the PR account. The most straightforward of these is that whenever a string is ambiguous between a PR and a RC, disambiguation toward a RC should come with a processing cost (PRFH). To explore this prediction, Pozniak et al. (2019) conducted an eye-tracking-while-reading study with native speakers of French (N = 69). For this, they manipulated two factors: the type of matrix predicate (perceptual vs non-perceptual) and its tense (match vs mismatch with the embedded predicate). As discussed above, PRs are licit after perceptual predicates but not after non-perceptual ones. As an additional restriction, we can simplistically say that the tensed element in a PR must match in tense with the matrix predicate.[38] This is demonstrated in (48) with a proper name to block a RC parse. No such restriction holds for RCs (49).

(48) a. *Gianni ha visto Maria che correva.* [Italian]
Gianni has seen.PST Maria that ran.PST

'Gianni saw Maria running.'

b. * *Gianni ha visto Maria che corre.*
Gianni has seen.PST Maria that runs.PRS

(49) *Gianni ha visto la donna che corre.* [Italian]
Gianni has seen.PST the woman that runs.PRS

'Gianni saw the woman that is running.'

[36] Here we only report values from their verbal condition which is comparable to the Italian results discussed above even though Grillo et al.'s (2015) experiment and statistical modelling also included nominal conditions (e.g., the sound of the grandma of the girl that was screaming). However the effect of *predicate* was not found to interface with *environment* (i.e., verbal or nominal). For discussion of the nominal results, please see the original paper.

[37] This appears to have been run over only the verbal condition of the English data, although this was not explicitly stated.

[38] For additional nuance that is irrelevant to the present discussion see Grillo & Moulton (2016)

Therefore in items like (50a, b), the PRFH would predict an initial PR bias because the matrix predicate is a verb of perception. Upon encountering the embedded predicate, however, a PR parse becomes incompatible with the string in (50b) due to the tense mismatch. Therefore, the account would predict the embedded predicate in (50b) to be processed more slowly than in (50a). No such difference is expected between the tense-match and mismatch items in (50c, d) because the matrix predicates are non-perceptual. As such, we would predict an interaction of *predicate type* and *tense match* in the eye-tracking measures.

(50) a. *Jean a vu la fille qui | poussait | la femme.* [French]
Jean has seen.PST the girl that pushed.PST the woman

'Jean saw the girl (that pushed/pushing) the woman.'

b. *Jean voit la fille qui | poussait | la femme.*
Jean sees.PRS the girl that pushed.PST the woman

'Jean sees the girl that pushed the woman.'

c. *Jean était marié à la fille qui | poussait | la femme.*
Jean was.PST married to the girl that pushed.PST the woman

'Jean was married to the girl that pushed the woman.'

d. *Jean est marié à la fille qui | poussait | la femme.*
Jean is.PRS married to the girl that pushed.PST the woman

'Jean is married to the girl that pushed the woman.'

Consistent with the prediction above, the authors observed evidence for an interaction of *predicate type* and *tense match* in regression path duration (or 'Go Past') times for the embedded predicate (indicated by the vertical bars in 50) as well as the proportion of regressions out of the region. Based on visual inspection of the plotted model estimates, they argue that this was driven by greater processing cost for tense mismatch items when the matrix predicate was perceptual but not when it was non-perceptual. However, this effect was only present in the first half of the experiment. Nonetheless, when the authors compared the French results to results from native English speakers (N = 50) with translational equivalents, they observed evidence of a three-way interaction with *language*. As independent modelling of English results did not reveal a significant interaction between *predicate type* and *tense match*, Pozniak et al. (2019) interpret the three-way interaction with *language* to indicate that the advantage for tense match under predicates of perception is restricted to French, the language that admits PRs. Thus, on the one hand, this study provides initial online evidence in support of the PRFH. On the other, however, the effect disappeared by the second half of the experiment, suggesting that it is subtle and can be obfuscated by adaptation to the experimental context. In this particular experiment, this may have been due to the highly reliable nature of tense (as well as the relatively low rate of fillers, only 48% of the French items). As the embedded predicates were always in a past tense form, upon encountering a matrix perceptual predicate in the present, it is plausible that participants could eventually anticipate that the embedded CP could only be parsed as a RC.

Building off this, Aguilar et al. (2021) explored how PR-firstness interacts with attachment in online parsing with a less reliable cue. For this, they recruited 42 native (Iberian) Spanish speakers. Following the PR account, they expected contrasting biases when PRs are, and are not, locally available. When PRs are available, items disambiguated toward LA are expected to be harder to process than items disambiguated toward HA. This is because PR parses are assumed to be preferred over RC parses (PRFH) and PRs are incompatible with LA in the relevant contexts. However, when PRs are unavailable and the string must be parsed as containing a RC, LA disambiguated items are predicted to be easier to process than HA items due to principles of locality (e.g., LATE CLOSURE, Frazier 1978).

For their eye-tracking-while-reading experiment, they created items as in (51). In these items, they manipulated gender marking on secondary depictive predicates within the embedded CPs to force attachment and the type of matrix predicate (perceptual vs non-perceptual) to affect PR availability. As the authors counterbalanced both the gender of NP_1 and the gender marking of the secondary depictive predicate, this study lacked a reliable cue like the one in Pozniak et al. (2019). Moreover, fillers also included a number of unambiguous PRs as in (52) which has been observed to mitigate adaptation effects (Fernandes, Alexiadou, Chow, Santi & Grillo 2018).

(51) a. *Juan vio al entrenador de la tenista que lloraba amargad(-o/-a) por la*
Juan saw DOM.the coach.M of the tennis.player.F that cried bitter(.M/.F) for the
derrota. [Spanish]
defeat

'Juan saw the coach of the tennis player (that wept/weeping) bitterly for the defeat.'

b. *Juan conoció al entrenador de la tenista que lloraba amargad(-o/-a) por la*
Juan knows DOM.the coach.M of the tennis.player.F that cried bitter(.M/.F) for the
derrota.
defeat

'Juan has met the coach of the tennis player that wept bitterly for the defeat.'

(52) *El técnico de laboratorio observó a Rosa que estaba escribiendo las fórmulas en la*
the technician of lab observed DOM Rosa that was writing the formulas on the
pizarra. [Spanish]
board.

'The lab technician observed Rosa writing the formulas on the board.'

At the disambiguating secondary predicate, the authors reported a significant interaction of *predicate type* and *attachment* in early processing measures (first fixation and gaze or 'first-pass' duration) which was driven by a LA bias in the non-perceptual items. No HA bias was observed in perceptual ones. In the later total time measure, however, this pattern reversed. *Post hoc*s indicated that in that case the interaction was driven by a HA bias in perceptual items with no LA bias in the non-perceptual ones. As such, this experiment provides two pieces of evidence in support of the PR account. First, despite previous studies repeatedly categorising Spanish as a HA language, once PRs are controlled for, the authors find evidence for an online LA bias with true RCs. Second, at least in total times, there was also online evidence for a HA bias when the matrix predicates were compatible with PRs. However, this later interaction was further modulated by the order of presentation with the HA bias only appearing in the first half of items. Thus, despite the study lacking a reliable cue for PR/RC disambiguation, PR-firstness appeared susceptible to adaptation as in Pozniak et al. (2019). The authors suggest that this may have been caused by the fact that RCs outnumbered PRs even after the inclusion of PR filler items.

More recently, Lee & De Santo (2022) conducted a self-paced reading study with native Italian speakers (N = 74) in which participants were presented sentences like the one in (53). In these items, the authors manipulated the type of matrix predicate (perceptual vs non-perceptual) and the number agreement on the embedded predicate to force attachment high or low. After each item, participants were asked a *who*-question targeting the subject of the embedded CP. Offline results were non-significant. Online results however indicated a significant interaction of *predicate type* with *attachment* on the window containing the embedded verb. This surfaced in the non-perceptual items as longer reading times on the embedded verb when the RC was

forced to attach high. In the perceptual items, the effect of attachment was non-significant.

(53) a. *Gianni vide il figlio dei medici che correv(-a/-ano) la maratona.* [Italian]
Gianni saw the son.SG of.the doctors.PL that run(.SG/.PL) the marathon.

'Gianni saw the son of the doctors (that ran/ running) the marathon.'

b. *Gianni viveva con il figlio dei medici che correv(-a/-ano) la maratona.*
Gianni lived with the son.SG of.the doctors.PL that run(.SG/.PL) the marathon.

'Gianni lived with the son of the doctors that ran the marathon.'

However, De Santo & Lee (2022) presented an additional analysis for their experiment in Lee & De Santo (2022). Although they did not observe a significant effect of attachment on the critical region of perceptual items, when they investigated the immediately following word (*la* in 53), they observed significantly longer reading times for perceptual items when they were LA disambiguated than when they were HA disambiguated. That is to say, they observed an online effect of PR availability, but only at the post-disambiguating window.

Thus in summary, whereas there is growing evidence that PR availability modulates interpretive biases for 'RC' attachment ambiguities, the available online results are somewhat mixed. On the one hand, there is evidence in support of the PRFH in simple contexts which abstract away from attachment. On the other hand, of the two experiments that have previously looked at PR-firstness and attachment, one reports a HA bias in PR-compatible contexts at the critical window (Aguilar et al. 2021) while the other found this effect only at their relatively short post-critical window (Lee & De Santo 2022). Moreover, several of those studies that have reported online evidence for PR-firstness indicate that the effect is subtle and susceptible to adaptation despite the putative universality of the PRFH (Fernandes et al. 2018, Pozniak et al. 2019, Aguilar et al. 2021).

2.4.3 Attrition and attachment biases

Circling back to attrition, there is some initial online and offline evidence from Spanish that these biases may attrite. In an initial study, Dussias (2003) conducted two experiments to explore RC attachment biases in late Spanish-English bilinguals of both L1s (L1-Spanish: N = 31; L1-English: N = 32). Both groups of bilingual speakers were highly proficient in both languages and had nontrivial immersion in their respective L2s (L1-English mean = 2 years; L1-Spanish mean = 7.5 years). Two monolingual control groups were also recruited (L1-Spanish: N = 14; L1-English: N = 19).

In the first experiment, participants were presented an offline questionnaire with critical items like in (54) to tap attachment preferences for both languages (months apart). Consistent with previous results, the Spanish monolinguals selected HA interpretations significantly more than LA ones (74% HA). Similarly, the English controls selected HA interpretations significantly less than LA ones (14% HA). Turning to the bilingual responses in English, the two groups patterned together and, similar to monolingual controls, displayed a preference for LA interpretations (L1-English: 28% HA; L1-Spanish: 22% HA). No significant group-level differences were observed. As for their Spanish responses, again the two groups pattern similarly. In this context however, they both differed from the monolingual baseline and exhibited a LA bias (L1-English: 44% HA; L1-Spanish: 28% HA; Monolingual Spanish: 74% HA). This difference was significant between the L1-Spanish bilinguals and the monolingual controls but not between the two bilingual groups. Thus, it appears the L1-Spanish speakers who are undergoing attrition exhibit an L2-like LA bias.

(54) a. Peter fell in love with the daughter of the psychologist who studied in California.

b. *Pedro se enamoró de la hija del psicólogo que estudió en California*[Spanish]
Pedro REFL make.love of the daughter of.the psychologist that studied in California

' Peter fell in love with the daughter of the psychologist who studied in California.'

Building on this in a second experiment, Dussias (2003) explored online parsing preference in Spanish via a (non-cumulative) self-paced reading task. In that task, items were forced toward high or low attachment via a PP in the embedded CP which biased attachment via world knowledge. This is exemplified in (55) with vertical bars to indicate window boundaries. Given that Dussias's (2003) work predated the PR account, they predicted a global HA bias for their Spanish monolinguals. That is to say, they predicted that reading times for the final PP should be slower when the embedded CP was biased to attach low, than when it was biased to attach high. As they formulated their predictions for the self-paced reading experiment prior to running their sentence interpretation task, the bilingual L1-Spanish speakers were not predicted to differ from the control group. Analysis of the reading times for the disambiguating PP indicated that the monolingual Spanish controls read HA-biased sentences significantly faster than LA-biased ones. For the L1-Spanish bilinguals, however, this pattern was significant in the opposite direction, i.e., LA disambiguated items were read more quickly. No significant difference between the two relevant conditions was observed within the L2-Spanish bilinguals. The results from the L1-Spanish bilinguals however corroborate the findings from Dussias's (2003) first experiment. Namely, L1 speakers of Spanish whose monolingual counterparts exhibit a HA bias appear to exhibit a LA bias under attrition.

(55) a. *El perro mordió al cuñado de la maestra | que vivió en Chile | con su*
the dog bit DOM.the brother.in.law.M of the teacher.F that lived in Chile with their
esposo. [Spanish]
husband

'The dog bit the brother-in-law of the teacher who lived in Chile with her husband.'

b. *El perro mordió a la cuñada del maestra | que vivió en Chile | con su*
the dog bit DOM the sister.in.law.F of.the teacher.M that lived in Chile with their
esposo.
husband

'The dog bit the sister-in-law of the teacher who lived in Chile with her husband.'

Convergent results from eye-tracking-while-reading tasks have also been reported. For an eye-tracking-while-reading task, Dussias (2004) recruited two groups of native L1-Spanish speakers: (i) a control group of monolingual speakers (N = 36) and (ii) a group of L2-English speakers (N = 20) who had been immersed in their L2 for an average of 3.7 years. To force attachment within their items, the NPs in the matrix CP always differed in gender such that only one could agree with the gender marking on the adjective in the embedded CP (*enferma* – 'sick' in 56). For their analysis, the author focused on the FP and TT eye-tracking measures on the disambiguating adjective (indicated by vertical bars). Pairwise comparisons indicated that the control group exhibited a significant HA bias in the TT measures whereas the experimental group exhibited a LA bias in both the FP and TT measures. As such, Dussias's (2004) eye-tracking results are fully consistent with the self-paced reading results in Dussias (2003).

(56) a. *El policía arrestó al hermando de la niñera que estaba | enferma | desde hacía*
the police arrested DOM.the brother.M of the baby.sitterF that was sickF since for
tiempo. [Spanish]
time

'The police arrested the brother of the babysitter that had been sick for a long time.'

b. *El policía arrestó a la hermanda del criado que estaba | enferma | desde hacía*
the police arrested DOM the sister.F of.the servantM that was | sickF | since for
tiempo.
time

'The police arrested the sister of the servant that had been sick for a long time.'

Building on this, Dussias & Sagarra (2007) conducted another eye-tracking while reading experiment with three groups of native (Iberian) Spanish speakers: (i) a control group of functionally monolingual speakers (N = 54), (ii) an experimental group (N = 28) with limited immersion in their L2-English (mean LoR = 8.5 months), (iii) and an experimental group (N = 20) with non-trivial immersion in their L2 (mean LoR = 7.1 years). Items were identical to those used in Dussias (2004). Analysis of the FP and TT measures for the disambiguating adjective indicated a significant interaction of *group* and *attachment* in the TT measures. *Post hoc*s indicated that the control group and the limited immersion groups exhibited a significant HA bias whereas the extensively immersed bilingual group exhibited a significant LA bias.

In additional modelling (which appears to have only considered the bilingual data) they also considered the potential effects of *LoR* and *L2-proficiency*. For the latter, two measures were considered: the same bilingual participants reading times for filler items in an unrelated experiment conducted in English, and their accuracy to comprehension questions for the same items. For the dependent variable, the authors calculated a 'disambiguation effect' by subtracting participants' mean reading times for HA items from LA ones (i.e., a positive value indicates a more monolingual-like HA bias). That model indicated that only *LoR* significantly affected the 'disambiguation effect,' with longer stays in the L2 associated with a greater shift toward a LA bias. The original authors take this to indicate that the group level differences observed between the two bilingual groups should be attributed to differences in the exposure to the L2 rather than any differences in their L2 proficiency.

Tying these studies together, it appears that the shift from a HA bias to a LA one in Spanish under attrition is a replicable effect that should be accounted for. However, it is unclear how we can integrate this with the PR account as all of the previous work on Spanish predated Grillo (2012) and therefore did not consider PR availability. Nonetheless, comparing these studies to those discussed in section 2.4.2 the offline and online results from native speakers undergoing attrition appear similar to monolingual speakers when PRs are locally blocked i.e. in both cases speakers exhibit a LA bias where they would otherwise be expected to exhibit a HA bias. As such, one might wonder whether attrition somehow affects the availability of PRs or their parsing. Relating this back to the work on aspects of null-subject languages, this could potentially relate to surface overlap with the L2 in which the relevant strings could only be parsed as RCs in turn leading to the observed LA bias. Although we did not have access to the stimuli from Dussias (2003), the fact that 37.5% of the stimuli from Dussias (2004) and Dussias & Sagarra (2007) contained PR-taking matrix predicates (Grillo & Costa 2014, fn. 29)[39] lends some plausibility to this hypothesised change to PR parsing. However, this has yet to be experimentally explored. Moreover, given that it is unclear how we can integrate these attrition results with the PR account, it is also unclear how these results fit into the larger discussion of attrition. It is against this background that the research questions for the present thesis were drawn.

2.5 Research questions

Against the general background above, we posed the very general starting question in (RQ_0). Recall that this question was motivated by the fact that previous accounts of attrition, which have focused on aspects of null-subject languages, have sought to attribute attrition to either a relaxation of featural specifications

[39] Dussias (2004) and Dussias & Sagarra (2007) used the items in Carreiras & Clifton (1999)

(e.g., Tsimpli et al. 2004) or overuse of a 'processing default' due to the cognitive complexity of reference maintenance (e.g., Sorace 2011). However, we have argued that many of the previous results relating to the null-subject parameter could alternatively be taken as evidence of a general tendency to rely more on interpretive biases already present in the L1. Attachment ambiguities offer a test case for this idea as we have argued that they do not sit at interface with discourse, but are instead conditioned by principles of computational efficiency. Therefore, if there is something about more external interface structures which causes them to be more susceptible to attrite, attrition should affect the two phenomena differently. If instead, attrition affects the two phenomena similarly, we should favour the 'increased biases' account for reasons of parsimony.

(RQ_0) Does attrition affect syntax-discourse structures and attachment ambiguities similarly?

To explore (RQ_0), the present thesis presents experiments targeting both offline interpretive biases and online processing of the two phenomena highlighted above: intrasentential pronominal resolution and PR/RC parsing ambiguities. For these experiments, a series of smaller, more manageable, research questions were formulated.

Starting with chapter 3, we investigate interpretive biases for both phenomena in globally ambiguous contexts. In that experiment, we are primarily interested in whether the effect observed in Dussias (2003) generalises to Italian, and if so how does the effect relate to PR-availability. Within the same experiment, we will also consider pronominal resolution using an adapted version of the task in Tsimpli et al. (2004) such that we can compare attrition effects within the same task and participants.

(RQ_1) Do (pseudo-)relative clause attachment biases in L1-Italian shift toward LA under immersion in L2-English as reported for L1-Spanish speakers?

(RQ_2) If so, is this change in attachment biases limited to environments in which Italian allows PRs?

(RQ_3) In the same participants, can we observe changes in the $\pm$ TS interpretations for null or overt pronouns in ambiguous items as reported for the speakers of various null-subject languages?

Despite chapters 5 and 6[40] presenting results from a single experimental session, pronominal and PR/RC results are presented in separate chapters for expositional simplicity. Within the pronominal items (chapter 5), we are primarily interested in investigating whether we can detect changes to participants' online sensitivity to the PAH given that only one of the studies reviewed above (i.e., Chamorro et al. 2015a) reported unambiguous evidence for attrition at the group level.

(RQ_4) Within native Italian speakers still living in Italy, can we observe both PAS effects in our intrasentential pronominal items?

i. The null pronoun prefers an antecedent which is in SpecIP.

ii. The overt pronoun prefers an antecedent which is not in SpecIP.

(RQ_5) If so, how are these biases affected by attrition?

In chapter 6, our goals are two-fold. First, given the relatively limited work on the role of PRs in the online parsing of attachment ambiguities, (RQ_6) and (RQ_7) are aimed at establishing the state of affairs in the Italian of native speakers still living in their L1 community. Second, we are interested in how sensitivity to the proposed attachment biases is affected by attrition such that we can compare the potential effects of attrition on the two phenomena of interest in the general discussion.

[40] Chapter 4 is devoted to the development and norming of the items for chapters 5 and 6, and as such does not have corresponding research questions.

(RQ_6) In native Italian speakers still living in Italy, can we observe a locality effect in the processing of items that unambiguously contain RCs?

(RQ_7) In native Italian speakers still living in Italy, can we observe a PR-firstness effect in the processing of items that are PR-compatible?

(RQ_8) If so, how are these biases affected by attrition?

2.6 Chapter summary

The present chapter introduced the two phenomena in which we will explore attrition throughout the rest of the thesis (pronominal resolution and PR/RC parsing ambiguities). Comparison of attrition in these structures was motivated by a review of the previous literature on attrition and aspects of null-subject languages. From that literature, we have argued for two general patterns under attrition: the often discussed overextension of certain discourse-conditioned forms (potentially due to surface-level overlap) as well as a general increased reliance on interpretive biases for other forms which cannot be directly attributed to transfer (be it at the level of representations or processing) from the speakers' L2s. Moreover, we have argued that neither of these patterns are fully compatible with accounts for attrition couched in terms of featural specifications or overt pronouns as a processing default. This led us to pose a series of research questions to explore the potential convergence between attrition at the syntax-discourse interface and beyond in order to investigate the role of interface type or whether we should pursue a more general bias-centric account of attrition.

3 Attrition and interpretive biases in globally ambiguous items

3.1 Introduction

In this chapter, we present a sentence interpretation task conducted to tap offline interpretive biases in ambiguous pronominal and PR/RC sentences. This was conducted with both Italians in their home community and Italians living abroad using items already available in the literature (section 3.3) for maximal comparability with previous work. The results from the present study (section 3.4) are interesting for two reasons. First, we demonstrate that the increased SpecIP bias for null pronouns in the Italian of speakers undergoing attrition reported in Tsimpli et al. (2004) is replicable. We take this as additional evidence that this pattern should not be treated as spurious but rather accounted for. Second, our results for the PR/RC items bear implications for how we understand the attrition of attachment biases. *Pace* Dussias (2003), we argue that attrition of attachment biases leads to an apparent strengthening of a (hidden) low attachment bias already present in the L1, rather than a loss of a global high attachment bias (section 3.5). We also discuss (dis-)similarities between our pronominal and PR/RC data.

3.2 Research questions

In this chapter, we focus on the following research questions:

(RQ_1) Do (pseudo-)relative clause attachment biases in L1-Italian shift toward LA under immersion in L2-English as reported for L1-Spanish speakers (Dussias 2003)?

(RQ_2) If so, is this change in attachment biases limited to environments in which Italian allows PRs?

(RQ_3) In the same participants, can we observe changes in the ± TS interpretations for null or overt pronouns in ambiguous items as reported for the speakers of various null-subject languages?

3.3 Method

3.3.1 Participants

Two groups of participants were collected: an experimental group and a control group. Regardless of group, all participants (i) were native Italian speakers, (ii) had lived in Italy from birth until at least the age of 16, (iii) had no diagnosed language-related disorders, and (iv) reported growing up monolingually.[41]

The control group consisted of 33 Italian native speakers (12 women, 21 men) who were living in Italy. On average, the control group was 40.91 years of age (SD = 6.40 years)[42] at the time of testing. As part of our background questionnaire, all participants were asked how many languages they spoke at or above an 'intermediate' level as well as for how many hours they used each of these languages in a typical day. This

[41] Although all of our participants reported growing up monolingually, when specifically asked about potential exposure to 'dialect' as a child (for example with one's grandparents) a portion of our sample indicated some exposure (Control: N = 5; Experimental: N = 5).

[42] Note, although our groups are somewhat older than those in studies such as Chamorro et al. (2015a) or Martín-Villena (2023), they are more comparable to those in Tsimpli et al. (2004: Ianthi Tsimpli, p.c.). Moreover, given that the ages between the experimental and control groups are matched, the age of our participants is unlikely to drive any potential group-level effects. This contrasts cleanly with Kaltsa et al.'s (2015) study in which they compared 2 groups of younger participants (mean age = 29 years) to 2 groups of older participants (mean age = 59 years, i.e., 30 years older than the younger group), and reported that the older groups (regardless of attrition status) behaved differently than the younger groups (regardless of heritage language status).

revealed that the control participants were not idealised monolinguals but to some extent multilingual. Of the 33 controls, 30 reported speaking some language other than Italian. Of those 30 participants, all reported speaking English. To quantify how much of each language a participant used, we calculated their percentages for each language (i.e., hours of language x / total hours reported for any language). Table 3.1 presents group means and standard deviations for the use of Italian, English, and 'Other' as percentages of a typical day.

	Control		Experimental	
	Mean	SD	Mean	SD
Italian	84.95%	13.67%	22.96%	18.69%
English	11.89%	11.78%	76.16%	18.53%
Other	3.17%	5.79%	0.87%	2.63%

Table 3.1: Mean language use by participants as percentages of a typical day with standard deviations.

The experimental group consisted of 29 Italian native speakers (18 women, 11 men) who were living in a majority English-speaking country at the time of testing.[43] Their minimum length of residency was 6 years (cf. Tsimpli et al. 2004), but on average this was considerably higher (mean = 14.27 years; SD = 7.89 years). We also pre-screened participants in the experimental group to ensure that none of them had travelled back to the L1 community in the three months prior to testing to avoid any possible re-exposure effects (Chamorro et al. 2015a). At the time of testing, the experimental group was 43.48 years of age (SD = 7.18 years).[42] As with the control group, we asked participants how many languages they spoke at or above an 'intermediate' level as well as for how many hours they used each of these languages in a typical day. We then calculated the percentages of use for each language. Group means and standard deviations for the use of Italian, English, and 'Other' as percentages of a typical day are presented in Table 3.1. Furthermore, given that some previous work on attrition had restricted itself to 'near-natives' (e.g. Tsimpli et al. 2004, Chamorro et al. 2015a), we also had the experimental group take the Cambridge Assessment General English quick placement test.[44] On average the experimental group scored 21.76 / 25 (SD = 1.66) suggesting that they are likely upper-intermediate to advanced L2 speakers. As such, following previous work, the experimental group's prolonged L2 immersion, lack of recent re-exposure to the L1 community, frequent L2 use, attenuated L1 exposure, and high L2 proficiency indicate that this is a suitable group in which to look for syntactic attrition.

3.3.2 Stimuli

For this experiment, items (N = 70) consisted of three types: pronominal (section 3.3.2.1), PR/RC (section 3.3.2.2), and filler items. Full lists of the two critical item types are available in Appendix B.

Similar to the PR/RC and pronominal items, the filler items (N = 26) were always biclausal. However, they were also always globally unambiguous. An example is provided in (57).

(57) *Giada gioca a volley mentre Irene gioca a basket.* [Italian]
Giada plays at volleyball while Irene plays at basketball

'Giada plays volleyball while Irene plays basketball.'

[43] Note, as all participants reported growing up monolingually in Italy, the experimental participants qualify as first-generation immigrants (i.e., potential attriters), not second-generation immigrants (i.e., heritage speakers).

[44] Items for the Cambridge Assessment General English quick placement test available at: https://www.cambridgeenglish.org/test-your-english/

These items were included for two reasons. First, they were included as genuine fillers to dilute the experimental items such that the PR/RC and pronominal items each only accounted for approximately a third of the total items (34.29% and 28.57% respectively). Second, these items served as an attentional control. Given that fillers were unambiguous, we could identify a correct answer for the compression questions that followed filler items. If a participant failed to achieve a sufficiently high accuracy rate on the filler items (implemented as 85% accuracy), they were deemed to have paid insufficient attention and were removed from the sample prior to the analysis of the two experimental item types.[45]

3.3.2.1 Pronominal items

Pronominal items (N = 20) were adapted from Tsimpli et al.'s (2004) picture selection task. These sentences vary in two dimensions: pronoun (null or overt) and direction (Matrix CP first vs Embedded CP first). As the original items do not appear to have been fully counterbalanced, we decided to adapt the items by creating two lists in which we counterbalanced the pronominal subject in each sentence (i.e., for an item like 58a, b which appears to have originally been presented only with an overt pronoun, we also created a version with a null pronoun.). This was done to ensure that any effect of pronominal type in later modelling was not due to differences in the items used in the null and overt conditions. We decided not to additionally counterbalance direction for two reasons. First, we are not theoretically interested in the effect of direction and as such, counterbalancing that manipulation was deemed to be neither particularly useful nor strictly necessary. Second, by not counterbalancing direction, we can compare our results to those in Tsimpli et al. (2004) within exactly the same items. As such, we take this partial counterbalancing to be an optimal (although not ideal) compromise between comparability and controlling for potential confounds. An example of the four conditions for the pronominal items is presented in (58) using two sentence pairs. These sentences were distributed across two lists such that each participant saw only one sentence per pair, 5 in each condition.

(58) a. Matrix-first[46] null condition [Italian]

L'anziana signora saluta la ragazza [quando ø attraversa la strada].
The.old woman greets the girl when *pro* crosses the road

'The old woman greets the girl when she crosses the road.'

b. Matrix-first overt condition

L'anziana signora saluta la ragazza [quando lei attraversa la strada].

'The old woman greets the girl when she crosses the road.'

c. Embedded-first[46] null condition

[Mentre ø guarda l'orologio] l'anziana signora si avvicina alla donna delle
While *pro* looks the.clock the.old woman REFL moves.closer to.the woman of.the
pulizie.
cleaning.

'While she looks at the clock the old woman moves toward the cleaning lady.'

d. Embedded-first overt condition

[Mentre lei guarda l'orologio] l'anziana signora si avvicina alla donna delle pulizie.

'While she looks at the clock the old woman moves toward the cleaning lady.'

[45] This resulted in the exclusion of two potential participants who were not included in the description above (section 3.3.1)

3.3.2.2 PR/RC items

The PR/RC items were borrowed directly from Grillo & Costa (2014). These items consist of 24 globally ambiguous sentence pairs as in (59). Within sentence pairs, only the matrix predicate varied. In the first condition, items contained verbs of perception (e.g. *vedere* - 'see' or *sentire* - 'hear') while in the second condition items contained non-perceptive predicates (e.g. *vivere con* - 'to live with' or *lavorare con* - 'works with'). Across all items, the matrix predicate was followed by a complex DP containing two NPs which in turn was followed by the embedded CP. The preposition between the two NPs was always *di* ('of') and the two NPs were always animate, singular, and appeared with the definite article. As with the pronominal items, sentences were distributed across two lists such that each participant saw only one sentence per pair, 12 in each condition.[47]

(59) a. Perceptual condition [Italian]

Gianni ha visto il figlio del medico che correva la maratona.
Gianni has seen the son of.the doctor that ran the marathon

'Gianni saw the son of the doctor running the marathon.'

b. Non-perceptual condition

Gianni vive con il figlio del medico che correva la maratona.
Gianni lives with the son of.the doctor that ran the marathon

'Gianni lives with the son of the doctor that was running the marathon.'

3.3.3 Procedure

Participants were recruited via Prolific (www.prolific.co) and paid £4.50 for their participation. The experiment was then run online through PCIbex (Zehr & Schwarz 2018). Participants first filled out a background language questionnaire and took the English placement test (for the experimental group). To explore interpretive biases, we then ran a sentence interpretation task. Participants were instructed that they would read a series of sentences in Italian and responds to related interpretation questions. Stimuli were presented alone on the participant's screen. When they finished reading the sentence, they pressed the space bar which caused the sentence to disappear and an interpretation question to take its place.[48] For PR/RC and pronominal items, the interpretation question always asked about the non-matrix CP. In filler items, the interpretation question asked about each CP (i.e. first or second) in half of items. An example of a PR/RC item with its corresponding question is presented in (60).

(60) *Gianni ha visto il figlio del medico che correva.* [Italian]

'Gianni saw the son of the doctor running.'

[46] Due to some confusion surrounding the terms 'forward' and 'backward' pronominal resolution, we will instead use the terms 'matrix-first' and 'embedded-first' to refer to which of the CPs appeared on the left. As such, our matrix-first items correspond to the backward items in Tsimpli et al. (2004), whereas our embedded-first items correspond to their forward items.

[47] The decision to re-use materials already present in the literature led to a noticeable imbalance in the number of items per condition in the two stimuli types (pronominal: 5; PR/RC: 12). Although this was undesirable, we decided that maximal comparability with previous works was preferable to a fully balanced design.

[48] For consistency with the PR/RC items, we did not use the images from Tsimpli et al.'s (2004) original picture selection task for the pronominal items.

Q. *Chi correva?*

'Who was running.'

F. *Il figlio*

'The son'

J. *Il medico*

'The doctor'

B. *Qualcun altro*

'Someone else'

Possible answers always consisted of the two relevant NPs from the sentential context as well as the option of *qualcun altro* ('someone else'). This third option was included as it is a reasonable answer for the pronominal items. The possible responses were counterbalanced by position such that NP_1 and NP_2 each appeared first in 50% of items and second in the other 50%. *Qualcun altro* ('someone else') always appeared last. The possible responses were labelled 'F,' 'J,' and ''B.' Participants could only select one interpretation. To do so, they pressed the relevant key on their keyboard. This caused the experiment to pass to the next item. Prior to data collection, ethical approval was obtained from the Ethics Committee of the Faculty of Modern and Medieval Languages and Linguistics at the University of Cambridge.

3.3.4 Data cleaning and analysis

Due to a coding error, we could not analyse two of the embedded-first pronominal items.[49] This affected one null and one overt item per participant. For the remaining pronominal and PR/RC data, we first calculated reading times for the entire sentences, as well as reaction time to the comprehension question. We then coded as missing any items with implausibly fast reading times ($<$ 1,500 ms) or decision times ($<$ 500 ms). This affected 0.54% of the data for the PR/RC items as well as 0.72% of the remaining pronominal items.

For the analysis of the pronominal items, we decided prior to data collection that we would analyse the matrix-first and embedded-first items separately. This was motivated by the fact that previous work has generally found different global patterns for the two directions. While it might be interesting to ask whether attrition significantly affects one of the directions more than the other, this was not an explicit question for the present thesis. Rather, we are specifically interested in how any of the pronoun-antecedent biases are affected.

Prior to conducting the experiment, we additionally decided to collapse the distinction between non-subject and 'other' readings. Although we included both as they are natural responses to the target questions and increased the comparability to Tsimpli et al. (2004), we are only theoretically interested in the [±TS] distinction. As such, the collapsing of the object and 'other' readings allowed us to more directly focus on the relevant feature and better compare attrition at the syntax-discourse interface and at the parser. As an additional benefit, this also allowed us to maintain a consistent use of logistic regressions between the two types of critical items.

To analyse both the matrix- and embedded-first pronominal data, responses (coded as ±SpecIP) were entered into a mixed effect logistic regression using the *lme4* package (Bates, Mächler, Bolker & Walker 2015) in R (R Core Team 2022). Using an empty model, we selected the best fitting random-effects structure for our within participant fixed effects (i.e., pronoun, Matuschek, Kliegl, Vasishth, Baayen & Bates 2017) using the Akaike Information Criterion (AIC). We then entered *pronoun* (negative level = null) and *group* (negative level = control) as fixed effects using sum coding (i.e., -0.5, 0.5).

[49] Items 16 and 18 in Table B.2

For the cleaning of PR/RC data, we additionally excluded 'someone else' responses. Recall that this option was included as a possible response as it was a natural response to our pronominal items. In the PR/RC items, however, it is nonsensical. Therefore, we took any such responses to those items to be accidental and coded them as missing. This affected a further 1.14% of the data for a total data loss of 1.68% in the PR/RC items.

For the modelling of the PR/RC data, *attachment* (coded as ±HA) was entered as the dependent variable. We then followed the same procedure as above to identify the best fitting random-effects structure before entering *predicate* (negative level: non-perceptual) and *group* (negative level: control) as fixed predictors using sum coding.

3.3.5 Hypotheses and predictions

For the pronominal items, we had the following hypotheses which were directly informed by the results reported in Tsimpli et al. (2004):

($H_{3.1}$) For overt pronouns in the matrix first items, the experimental group will select more non-TS (i.e. subject) interpretations than the control group.

($H_{3.2}$) For null pronouns in the matrix first items, the experimental group will select more non-TS interpretations than the control group.

($H_{3.3}$) In the embedded-first items, the experimental and control groups will not differ with regard to ± TS readings for null or overt pronouns.

Therefore, within the matrix-first items we expected a main effect of *group* as well as a main effect of *pronoun*. Specifically, for null pronouns in that direction, we expected the control group to select subject interpretations around chance level while the experimental group would display a SpecIP-bias. For overt pronouns in the same direction, both groups were expected to exhibit a non-SpecIP bias, although this is expected to be weaker in the experimental group.

In the embedded-first direction, we only expected an effect of *pronoun* with both the control and experimental groups exhibiting a non-TS bias for null pronouns and a TS bias for overt pronouns. As we have restricted ourselves to ± TS readings, we did not expect an effect of attrition with overt pronouns as reported in Tsimpli et al. (2004). This is because their effect was driven by a change in the rate of 'other' responses (as opposed to subject or object readings) rather than a change in TS readings. which is the focus of the present analysis.

For the PR/RC items, the following hypotheses were drawn:

($H_{3.4}$) The experimental group will exhibit a weaker global HA bias than the control group (cf. Dussias 2003) Therefore, we expect a main effect of Group.

($H_{3.5}$) This weakening of the global HA bias will be driven by a change in the perceptual items only.

These hypotheses were empirically motivated by the observation that offline and online results from native speakers undergoing attrition pattern similarly to monolingual speakers when PRs are locally blocked. As such, we expect an interaction of *group* by *predicate*. This is expected to surface as an effect of *group* within PR/RC items. Specifically, we expect the control group, but not the experimental group, to exhibit a HA bias. No effect of *group* is expected within the RC-only items. In that context, both groups are expected to exhibit a LA bias.

3.4 Results

3.4.1 Pronominal items

Figure 3.1 presents the percentage of subject responses broken down by direction, group, and pronominal form. In the embedded-first items,[50] the control and experimental groups both preferentially selected subject interpretations for null pronouns (control: 85.60%; experimental: 87.93%). For overt pronouns, they again patterned together in preferring non-subject interpretations although this was slightly more pronounced for the control group (control: 22.73%; experimental: 31.03%). In the matrix-first items,[51] the group differences were somewhat more noticeable. Whereas the control group selected subject interpretations for null pronoun at around chance levels, there was a slight subject bias in the experimental group (control: 46.97%; experimental: 62.22%). In the overt items, both groups preferred non-subject interpretations (control: 12.88%; experimental: 8.28%). The model outputs for the logistic regressions over the matrix- and embedded-first pronominal items are presented in Table 3.2.

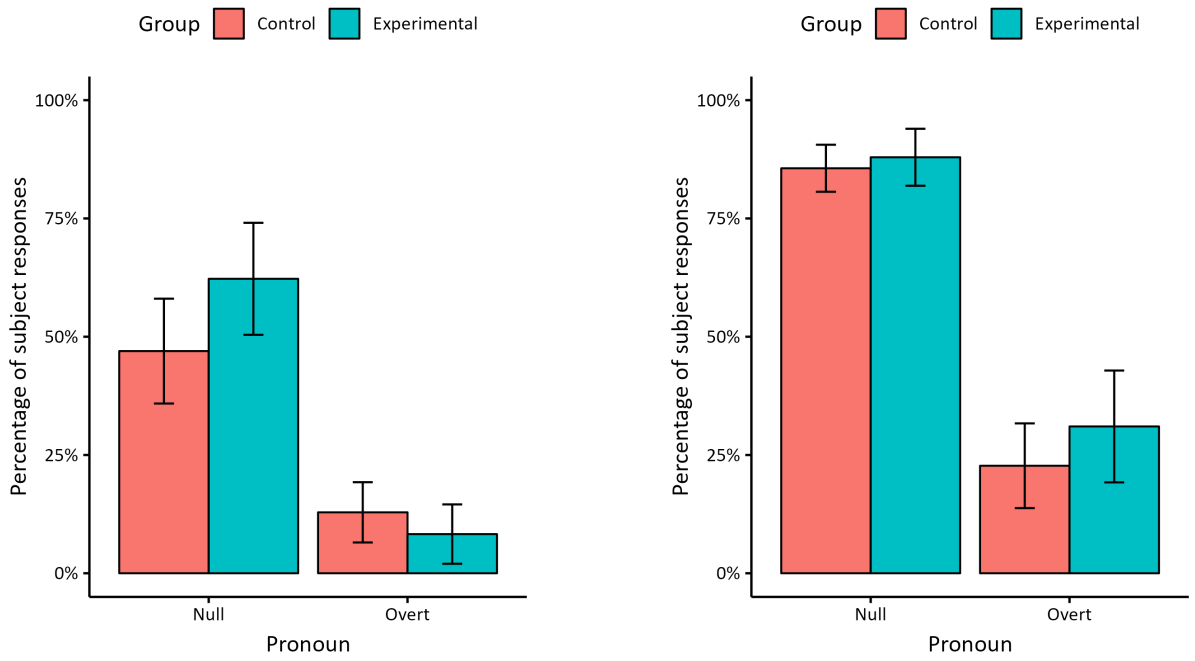

(a) Matrix-first.

(b) Embedded-first.

Figure 3.1: Subject response rates for the pronominal items by group, pronominal form, and direction with 95% confidence intervals.

The model for the matrix-first items indicated a significant effect of *pronoun* ($\hat{\beta}$ = -2.90; z = -3.91; $p < 0.001$). This surfaced as fewer subject interpretations with overt pronouns. The model did not indicate a significant effect of *group*. However the interaction between *pronoun* and *group* was significant ($\hat{\beta}$ = -1.47; z = -2.11; $p = 0.04$). The model for the embedded-first items only indicated a significant effect of *pronoun* ($\hat{\beta}$ = -3.46; z = -10.19; $p < 0.001$). This surfaced as fewer subject readings with overt pronouns. The effect of *group* and its interaction with *pronoun* were non-significant. To follow up on that significant interaction in the matrix-first items, we then conducted pairwise comparisons using the *emmeans* package (Lenth 2022). These comparisons with Holm–Bonferroni corrected p-values are presented in Table 3.3.

The pairwise comparisons indicated that the effect of *pronoun* was significant in both groups (control: $\hat{\beta}$

[50] For a breakdown of the two different types of non-SpecIP responses, please see Figure D.2.
[51] For a breakdown of the two different types of non-SpecIP responses, please see Figure D.1.

	Est.	Std. Error	*z*-value	*p*-value	
Intercept	-1.53	0.35	-4.43	<**0.001**	***
Pronoun	-2.90	0.74	-3.91	<**0.001**	***
Group	0.03	0.36	0.08	0.94	
Pro.:Group	-1.47	0.71	-2.11	**0.04**	*

(a) Matrix-first

	Est.	Std. Error	*z*-value	*p*-value	
Intercept	0.49	0.26	1.90	0.06	·
Pronoun	-3.46	0.34	-10.19	<**0.001**	***
Group	0.34	0.29	1.17	0.24	
Pro:Group	0.21	0.50	0.41	0.68	

(b) Embedded-first

Table 3.2: Model outputs for matrix- and embedded-first pronominal sentence interpretations.

	Estimate	Std. Error	*z*-value	*p*-value	
Pronoun within group					
Control	-2.16	0.75	-2.88	**0.01**	*
Experimental	-3.64	0.89	-4.11	<**0.001**	***
Group within pronoun					
Null	0.77	0.37	2.10	0.07	·
Overt	-0.72	0.62	-1.16	0.25	

Table 3.3: Pairwise comparisons for matrix-first pronominal interpretations with Holm–Bonferroni corrected *p*-values.

= -1.55; z = -4.89; $p < 0.001$; experimental: $\hat{\beta}$ = -2.85; z = -7.13; $p < 0.001$). In both cases, this was due to fewer subject interpretations with overt pronouns. The effect of *group* within the overt pronominal items was non-significant, and only marginal within the null pronominal items ($\hat{\beta} = 0.71$; $z = 2.17$; $p = 0.07$).

3.4.2 PR/RC items

Figure 3.2 presents the percentage of high attachment responses broken down by group and condition. In the PR/RC condition, both groups preferred high attachment at similar rates (Experimental: 66.67%; Control: 72.05%). In the RC-Only condition, both groups dispreferred high attachment, but this effect was slightly more pronounced in the experimental group (Experimental: 15.99%; Control: 29.52%). The model output for the logistic regression over these items is presented in Table 3.4.

	Estimate	Std. Error	*z*-value	*p*-value	
Intercept	-0.44	0.27	-1.63	0.10	
Predicate	3.23	0.33	9.67	<**0.001**	***
Group	-0.77	0.39	-1.97	**0.049**	*
Predicate:Group	0.88	0.50	1.76	0.08	·

Table 3.4: Model output for PR/RC sentence interpretations.

This model indicated a main effect of *predicate* ($\hat{\beta} = 3.23$; $z = 9.67$; $p < 0.001$). This surfaced as a greater rate of high attachment in the perceptual condition. The model also indicated a main effect of *Group*

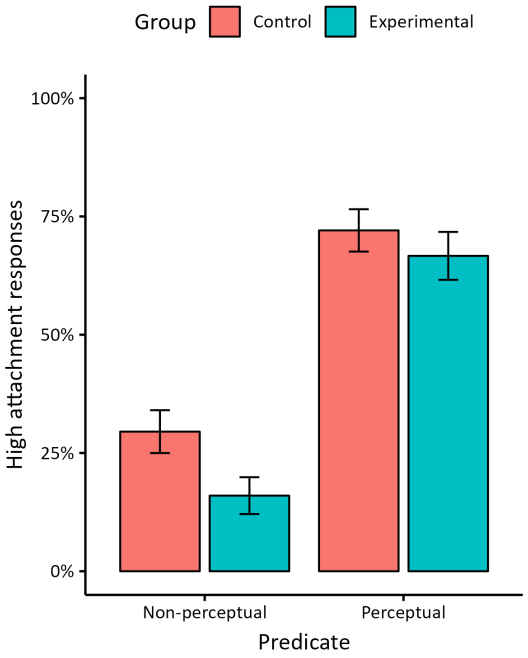

Figure 3.2: High attachment rates for the PR/RC items by group and predicate with 95% confidence intervals.

($\hat{\beta}$ = -0.77; z = -1.97; p = 0.049). This surfaced as a greater rate of high attachment in the control group (control: 50.70%; experimental: 41.03%). The interaction between *predicate* and *group* was not significant.

3.5 Discussion

3.5.1 Pronominal items

For the pronominal half of this experiment, we were interested in whether we could observe changes in ± TS interpretations as reported for the Italian participants in Tsimpli et al. (2004). Specifically, within the matrix-first items, we were looking for a shift toward more SpecIP interpretations with overt ($H_{3.1}$) as well as null pronouns ($H_{3.2}$) in the experimental group. In the embedded-first items, we only expected an effect of pronoun but no attrition ($H_{3.3}$).

Starting with the embedded-first items, results were remarkably similar to those in Tsimpli et al. (2004) with regard to SpecIP responses. Values for the null condition are within a few percentage points (their study: control: 85%;[52] experimental: 87%;[52]; our study: control: 85.60%; experimental: 87.93%), while values for the overt condition were slightly higher in general (their study: control: 13%;[52] experimental: 25%;[52]; our study: control: 22.73%; experimental: 31.03%). As such, we unsurprisingly observed an effect of *pronoun*. As predicted, this surfaced as a preference for non-SpecIP interpretation with overt pronouns as has been previously reported for native Italian speakers. This trend was unaffected by attrition in our analysis confirming ($H_{3.3}$). This result is thus convergent with those in Tsimpli et al. (2004) given that their attrition effect in these items was driven by a change in the rate of 'other' responses, rather than a change in [±TS] readings as was the focus of the present analysis (although see Figure D.2 for some visual evidence that that pattern likely was not replicated in our data.).

Turning to the matrix-first items, again our results were numerically quite similar to those in Tsimpli et al. (2004) with the control group selecting SpecIP readings around half the time with null pronouns while almost categorically rejecting them with overt pronouns (their study: null: 50.75%; overt: 7.6%; our study: null:

[52] Exact values for these conditions were not provided. As such, these are estimated from their Figures 4 and 5.

46.97%; overt: 12.88%). Looking at the experimental groups however, although our experimental group pattern with Tsimpli et al.'s (2004) in selecting more SpecIP readings that the control group (their study: 69.8%; our study: 62.22%), our experimental group selected fewer – not more – SpecIP readings for overt pronouns than our control group did (their study: 21.1%; our study: 8.28%). Statistical modelling over our results indicated a significant interaction between *group* and *pronoun*. Given that the *post hoc* tests within the two conditions were not significant after correcting our *p*-values, there seem to be two possible interpretations of this interaction. On the one hand, we might suggest that this interaction was driven by a change in the interpretation of null pronouns. Some support for this idea can be drawn from the fact that the non-SpecIP bias for overt pronouns was only larger by about 5 percentage points in the experimental group whereas the experimental group's the SpecIP bias for null pronouns was about 15 percentage points larger than the control group's. Moreover, the shift in the null items appeared in the predicted direction and matched the results of Tsimpli et al. (2004). On the other hand, we might suggest that this effect was driven by the numerically more exaggerated biases for both null and overt pronouns, despite the trend being noticeably smaller in the latter. Regardless of which interpretation the reader takes to be more reasonable, the implication of the present results remains the same. That is to say, the increased SpecIP interpretation for null pronouns in Italian under attrition is a replicable result confirming ($H_{3.2}$).

This raises the question of how we should interpret this attrition effect, given that our control group appear to have interpreted null pronouns as coreferential with the possible antecedent in SpecIP at around chance levels. We believe that this should be interpreted as an increased reliance on the SpecIP bias already present in the L1, which is to some degree obfuscated in these particular items. This is because substantial amounts of previous work using various sets of items in various intrasentential contexts have found native (adult) Italian speakers to exhibit a reliable SpecIP bias in both their interpretation and their processing of null pronouns (e.g., Carminati 2002, 2005, Filiaci et al. 2014, Vogelzang, Foppolo, Guasti, van Rijn & Hendriks 2019, Contemori & Di Domenico 2021). Conversely, previous work with these specific (matrix-first) items have found Italian control groups' SpecIP interpretation of null pronouns to hover around 50% (e.g., Tsimpli et al. 2004, Belletti et al. 2007), suggesting that they may contain some artefacts that slightly dampen the SpecIP bias for null pronouns. With this interpretation in mind, the fact that this effect in Italian is replicable, and the fact that a similar pattern is also replicable in Turkish (Gürel 2004, Gürel & Yılmaz 2011), we take this to mean that any successful theory of attrition should not treat changes to null pronouns as spurious but rather account for this fact.

As a final point on the pronominal results, it is striking that although the present study was able to replicate the infrequently-discussed attrition effect with null pronouns, we were unable to replicate their more frequently discussed relaxation of the non-SpecIP bias for overt pronouns *contra* ($H_{3.1}$). Neither do we have direct evidence for an increased non-SpecIP bias as was reported for Spanish speakers undergoing attrition by Martín-Villena (2023) who used very similar items.[53] This may relate to the fact that two tendencies observed and argued for in chapter 2 might compete in overt pronouns(i.e., relaxation due to overlap vs the general increased reliance on biases), resulting in the mixed pattern. However, we also cannot exclude the possibility that the lack of a significant effect within the overt items may relate to the fact that the control group in the present study (as well as the Italian control group in Tsimpli et al. 2004) were already approaching the ceiling for over pronouns in the matrix-first items (12.88% SpecIP), leaving very little room in which an increases non-SpecIP bias could manifest. We will return to this idea in chapter 5.

3.5.2 PR/RC items

For the PR/RC half of this experiment, we wanted to explore two questions. First, whether we could observe attrition of 'RC' attachment biases in a new language pair, L1-Italian L2-English. Second, whether attrition

[53] Recall that Martín-Villena (2023) manipulated the conjunctions of the original items from Tsimpli et al. (2004).

of these attachment biases is driven by a change in the parsing of pseudorelatives.

Starting with the first question, the results showed evidence of attrition. The experimental group selected HA interpretations significantly less frequently than the control group. Moreover, this effect was in the same direction as the attrition effect observed in Dussias (2003) confirming ($H_{3.4}$). Given that attrition of attachment biases may be observed in Italian as well as Spanish, we take this to suggest that the effect may be a general phenomenon for HA languages, at least under immersion in a LA-L2 like English.

Nonetheless, it is worth pointing out that the effect in the present study was less pronounced than the one reported in Dussias (2003). While Dussias observed a shift from 74% HA to 28% HA, we observed a global shift from 50.70% HA to 41.03% HA. An immediate question might be whether this is due to our participant sample. However as noted in section 3.3.1, our participants' prolonged L2 immersion, lack of re-exposure to the L1, frequent use of the L2, attenuated use of the L1, and high L2 proficiency all suggest this is an ideal group in which to look for attrition effects. Moreover, the present experimental group is patently undergoing attrition as attrition effects were attested in the pronominal items. Nonetheless, one might try to argue that our group is somehow less attrited than that in Dussias (2003). However, while it is hard to directly compare many of our background measures such as L2 proficiency with those in Dussias (2003), we can compare the length of L2 immersion, in which case our sample has been immersed for longer. For our participants, the average immersion in their L2 (mean = 14.27 years, SD = 7.89 years) was almost twice that reported in Dussias (2003: mean = 7.5 years). This is relevant because previous work on attachment ambiguities has found that length of L2 immersion is a significant predictor of L1 attrition within advanced L2 speakers (Dussias & Sagarra 2007). Thus, based on the available information, it is unlikely that differences in the size of the attrition effect can be attributed to our participant sample. Prior to exploring an alternative explanation, let us first bring PRs into the discussion.

Integrating PRs, our results indicated that the control group selected HA interpretations significantly more frequently when PRs were locally available (PR/RC: 72.05% HA; RC-only: 29.52% HA). Our values were in line with those reported in Grillo & Costa (2014: PR/RC: 78% HA; RC-only: 24% HA) and indicate that we can be confident the manipulation worked as intended.

Considering attrition in light of the PR/RC distinction then, we predicted that the attrition effect would be driven by a change in the PR/RC items ($H_{3.5}$). Recall that this was empirically motivated; the online and offline shift toward LA previously reported to occur under attrition was similar to what has been reported for monolingual speakers when PRs are locally blocked. Instead, we observed a global effect of *group* and *predicate* with no interaction between the two. This was unexpected and worth highlighting. Namely, the experimental group still exhibited a HA bias (66.67%) in PR-compatible contexts despite (i) a clear LA bias when PRs were blocked and (ii) a global shift toward LA. This suggests that the syntactic availability of PRs and their preferred parsing status over RCs are stable under attrition, a finding that is consistent with Grillo & Costa's (2014) idea that the PRFH is a universal parser bias. However, it raises the question of how to interpret the global shift toward LA responses in the experimental group if it is not attributable to a change in PRs.

As there are two theoretically relevant parser biases at play in our items, a reasonable place to start would be the other bias, locality. If there were an increased locality bias acting on true RCs, this would straightforwardly affect our RC-only items. However, it could have also reasonably affected our PR/RC items. To see how this would be the case requires going back to the responses from the control group and underscoring that biases are not categorical. In the PR/RC items, the control group selected HA in only 72.05% of trials i.e., they selected LA in the other 27.95%. As PRs are incompatible with LA in our items, this suggests that they parsed at least 27.95% of those trials as containing true RCs. Moreover, this is a lower bound estimate as even when items could only be parsed as RCs (i.e., RC-only items), the control group still selected HA interpretations in 29.52% of trials. Taken together this suggests that a non-trivial portion of the PR/RC items were parsed as containing RCs despite the PR bias. Those RC parses in the PR/RC condition would then have also been

susceptible to any increased locality effect in the same way as the items in the RC-only condition. As such, it seems reasonable to interpret our attrition effect as an increased locality effect within the experimental group. Moreover, despite the non-significance of the interaction term, it is worth pointing out that the shift toward LA within the RC-only items was roughly twice the size of the shift in the PR/RC items (13 percentage points vs 6). That is to say, the shift toward LA was numerically larger in the condition with more RC parses. As such, we take our data to be more consistent with this 'increased locality' interpretation than with our original PR-driven hypothesis ($H_{3.5}$) which would have predicted the opposite trend. Thus, the present results lead us to the opposite conclusion of that presented in Dussias (2003).

With the PR/RC distinction in mind, we can circle back to the observation that the attrition effect was smaller in the present study than in Dussias (2003). We suggest that this may be attributable, at least partially, to the characteristics of the experimental materials. Previous research on RC attachment ambiguities has observed that the length of the RC affects attachment. Specifically, longer RCs are more likely to result in HA (Same Size Sister Constraint, Fodor 1998). However, the items we borrowed from Grillo & Costa (2014) have very short RCs. In 22 of the 24 items, the RCs were maximally minimal, consisting of 2 orthographic words: the complementizer che ('that') and the tensed verb. In the remaining 2 items, one RC contained 3 orthographic words, and one contained 4. This likely increased the baseline LA bias for true RCs (29.52% HA for our monolinguals). This coupled with the fact that attrition appeared as an enhanced LA bias with true RCs, may have left little room for the attrition effect to manifest.

A further question raised by our results is how to account for the more pronounced locality effect in our experimental group. We can approach thinking about this apparent strengthening in two ways. Either the baseline LA bias was changed, or the participants' reliance thereon was changed. Let us explore these in turn.

As the speakers have another language which exhibits a LA bias for RCs, we might look for a potential explanation for the change in the L2. On the one hand, this is theoretically undesirable as it would require suggesting that the LA bias with RCs is somehow stronger in English than it is in Italian, i.e., this would force us to reject that there is a universal LA bias for RCs. On the other hand, if the LA bias for RCs differed in strength between the L1 and L2 and attachment biases can be modified experimentally as suggested by the Tuning Hypothesis (Mitchell & Cuetos, 1991, et seq.), this could potentially explain why RC parsing biases were affected while PR biases were not. Empirically, we cannot exclude this possibility. To the best of our knowledge, the only study that might bear on this issue is Grillo et al. (2015). As noted above, they tested L1 English speakers with direct translations of the stimuli in Grillo & Costa (2014). Their results indicated that their participants opted for HA in 19.5% and 21.1% of items in the RC-only condition. That puts their English results numerically below the monolingual groups in Grillo & Costa (2014, 24% HA) and the present study (29.52% HA) although still above our experimental group (15.99% HA). Moreover, when Grillo et al. (2015) statistically compared their English results with the Italian results from Grillo & Costa (2014), *post hoc*s only indicated a significant effect of *language* within the perceptual items. Within the non-perceptual items, native English and Italian speakers were not found to differ in the strength of the LA bias. Thus, on the one hand, we cannot empirically refute the idea that LA bias in the L2 affected the LA bias already present in the L1. However, on the other hand, the available evidence provides no indication that the two languages differ in the relevant bias, and to suggest that they do would be theoretically undesirable.

As an alternative to the above, we might suggest that the underlying LA bias remained unchanged but that the experimental group's reliance on structural biases may have been affected instead. Such an interpretation has two advantages. First, it allows us to maintain the universality of parsing. Second, it is consistent with our interpretation of the change in the null pronominal items. To maintain an account along these lines, however, we would require an explanation for why the LA bias with RC-Only items is affected but not the HA bias with PR/RC items. For this, we see at least two possibilities. First, the asymmetry could potentially relate to the type of parser bias implicated, i.e. the PRFH is about what structure is projected, while locality is about where an element is attached. Relatedly, the present study relied on the semantics of the matrix to

manipulate PR availability (i.e. only the verbs of perception could take events as their complements). This is relevant as it has been argued that the interface between syntax and semantics is relatively stable under attrition Sorace (2011). As such, our manipulation may have protected PR parsing. Future work may instead wish to focus on PRs in positions that are not selected by the matrix predicate (e.g., PRs may occupy matrix SpecIP in Italian) to better isolate their syntactic parsing bias.

Alternatively, the contrasting parser biases in the PR/RC items might have neutralised each other to some extent. Even if attrition also strengthened PR firstness resulting in more HA, we would not expect a bias to apply categorically. Therefore in a subset of cases, PR/RC items would likely still be parsed as containing RCs. These RC parses would then be subject to the enhanced locality effect noted above resulting in more LA. Ultimately, teasing apart these potential explanations is not possible with the available data.

Nonetheless, the fact that attrition affects the LA bias with true RCs which is shared between the L1 and L2 raises the question of whether we should generalise further. That is to say, if the conflict between global parser biases is only apparent as suggested by the pseudorelative account, can we observe a similar attrition effect between two languages which have the same 'global attachment strategy' (e.g. two LA or two HA languages), or even with other attachment ambiguities such as adverbials/PPs which are subject to a LA bias? To the best of our knowledge, the only study that might bear on attrition of RC attachment biases in a LA language is Dussias (2003). While she reports no significant difference between her controls and her L1-English speakers immersed in their L2-Spanish, the immersion period for that experimental group was relatively short (mean = 2 years). As such, they may not have been exposed for long enough for attrition effects to be observed (Dussias & Sagarra 2007). Therefore, the question of generalising attrition in biases beyond cases of apparent parser bias conflict should be explored in future work.

3.6 Chapter summary

In this chapter, we presented a sentence interpretation task containing globally ambiguous PR/RC and pronominal sentences. This was done to tap the offline interpretive biases for these types of ambiguities in two samples of native Italian speakers: one still living in Italy and one who had lived abroad for a number of years. Results indicated attrition effects for both phenomena in our sample. Contrary to our predictions for the PR/RC items, however, attrition for attachment biases was driven by an apparent strengthening of the LA bias for items with unambiguous RCs, not a loss of the HA bias when PRs were locally available. This result, although unexpected, was similar to the pattern observed for the pronominal items. There, we failed to find a relaxing of the non-SpecIP bias for overt items as reported by Tsimpli et al. (2004). However, we did observe an apparent strengthening of pronominal biases in the matrix-first items in line with their results for null pronouns.

4 Construction and evaluation of the locally ambiguous stimuli

4.1 Introduction

The present chapter details the construction and norming of the stimuli employed in the online experiments described in chapters 5 and 6. Items for these online experiments were custom-made for two reasons. First, while there are previous stimuli lists that have been used in experiments exploring pronominal resolution biases in Italian (e.g., Carminati 2002), the items used in previous studies on the online processing of relative clause (RC) attachment biases were either in other languages (e.g., Spanish: Aguilar et al. 2021) or did not manipulate pseudorelative (PR) availability (e.g., De Vincenzi & Job 1993). Second, in order to facilitate comparability between pronominal and PR/RC results, we wanted the two types of items to be maximally similar. We, therefore, constructed and evaluated two sets of new materials suitable for online measures guided by several desiderata:

(α) The stimuli must be somehow disambiguated.

(β) The availability of PRs should be directly manipulated within the PR/RC items.

(γ) The different versions of each item should be maximally similar with regard to non-structural factors (e.g. plausibility or phonological weight).

(δ) The disambiguating manipulation should be uniform across the two sets of items (PR/RC and pronominal).

(ε) The items should be suitable for both a self-paced reading and an eye-tracking-while-reading task.

The first of these desiderata (α) was motivated by the need to compare various readings of the same item (e.g. HA RC vs LA RC) in order to infer parsing biases. The second (β) was motivated by the string identity between PRs and RCs in the kinds of items explored. This necessitated some way of identifying when participants could parse items as PRs and when they could not (other than the end attachment interpretation). The third (γ) was included in order to allow us to focus on syntactic parsing biases to the exclusion of other factors that have been previously been demonstrated to affect parsing biases (e.g., phonological weight, Fernández 2003). The fourth (δ) was added in order to facilitate comparisons between the two phenomena (PR/RC attachment and pronominal resolution). The fifth (ϵ) was included for comparability across the separate experiments presented in the following chapters.

In the present chapter, section 4.2 details the development of the potential experimental materials to satisfy (α - ϵ). Section 4.3 presents how the materials were evaluated in order to: (i) ensure that the relevant readings were available and (ii) minimise any semantic/pragmatic biases within the final experimental items (γ). The final set of items derived from this process are reported in Appendix C and was identical across the self-paced reading and eye-tracking experiment described in the following chapters.

4.2 Development of the experimental materials

Starting with condition structure, we were interested in manipulating two binary factors in both the pronominal and PR/RC items. As such, there were four desired conditions for both types of critical items. In the PR/RC items we manipulated whether PRs were locally available (*predicate)* and to which of the possible NPs the

embedded CP attached (*attachment)*. In the pronominal items we manipulated the pronominal form of the subject in the embedded CP (*pronoun*) and whether it was co-referent with the matrix subject or object (*antecedent*). For the pronominal items, we did not manipulate the order of matrix and embedded clauses (*direction*) as in the sentence interpretation task (chapter 3). Instead, the matrix clause always proceeded the embedded clause containing the pronominal form. Second, The conditions for PR/RC and Pronominal items are presented in (61, 62) respectively.

(61) PR/RC conditions

a. Non-perceptual forced High

b. Non-perceptual forced Low

c. Perceptual forced High

d. Perceptual forced Low

(62) Pronominal conditions (matrix-first)

a. Overt pronoun with SpecIP reading

b. Overt pronoun with Non-SpecIP reading

c. Null pronoun with SpecIP reading

d. Null pronoun with Non-SpecIP reading

4.2.1 PR/RC items

For the PR/RC items, we started with the same base template (63) as in chapter 3.

(63) NP_0 - Matrix Pred. - NP_1 - *di* ('of') - NP_2 - Embedded CP

All NPs were animate. For consistency, NP_0 always consisted of a proper name (e.g., Gianni) while NP_1 and NP_2 were non-proper. In order to form felicitous complex NPs (i.e. NP_1 *di* NP_2), we used kinship terms for NP_1 (e.g., *cousin, aunt/uncle, son/daughter*). To side-step issues relating to contrast sets, we avoided using kinship terms (NP_1s) for which there may be a uniqueness presupposition (e.g., *marito* - 'husband'). This is because restrictive RC readings of the type that we are interested in carry with them a presupposition of a contrast set. To exemplify this, take (64). For a felicitous RC parse of the embedded CP, (64) presupposes the existence of at least one other (less happy) boy who was not eating gelato. In a context like (63) then, if NP_1 were to carry a presupposition of uniqueness, this may bias the RC to not attach high. In and of itself, this would create issues in the interpretation of any potential LA bias with RCs. This is because we would not be able to disentangle whether the effect were driven by (i) the locality bias we are interested in or (ii) a bias to avoid violating presuppositions regarding uniqueness/contrast sets. Integrating PRs however, issues compound. Unlike RCs, PRs do not carry a presupposition of a contrast set. As such, if we were to include NP_1 such as *marito* ('husband'), then if we were to observe a greater rate of HA in the PR/RC items than in the RC ones, we would not be able to disentangle whether the effect were due to the PRFH or the infelicity of HA RC interpretations.

(64) *Gianni ha visto il ragazzo che mangiava il gelato.* [Italian]
Gianni has seen the boy that ate the gelato

'Gianni saw the boy (that was) eating gelato.'

To force attachment within the PR/RC items (α), we settled on a gender manipulation. For this reason, NP_1 and NP_2 always unambiguously differed in gender. The grammatical gender of an NP was always recoverable either from the final vowel on the NP (e.g., $-o_M$, $-a_F$) and/or the NP's associated article. For consistency, we avoided NPs which are lexically specified for gender in Italian (e.g., *fratello* - 'brother', *sorella* - 'sister'). To avoid complicating the list structure, the gender of NP_1 was counterbalanced *across* – not *within* – items. Within the embedded CP we then included a secondary depictive predicate (e.g., *ubriaco* - 'drunk'). In Italian these elements overtly manifest gender agreement via a suffix ($-o_M$, $-a_F$). As the gender marking on the secondary predicate must agree with the subject of the embedded CP, and the subject of the embedded CP must match the gender features of the NP it attaches to, this allowed us to force attachment high or low. The modified template for the PR/RC items is sketched in (65).

(65) NP_0 - Mat. Pred. - $NP_{1,\ F}$ - *di* ('of') - $NP_{2,\ M}$ - *che* ('that') - Emb. Pred. - Secondary $Pred._{M/F}$ - ...

We opted for a gender manipulation over a number manipulation for two reasons. First, we wanted to control phonological weight (γ). Number - but not gender - is obligatorily marked on tensed elements in Italian. In the third person, singular (indicative) agreement surfaces as a monosyllabic suffix (-a, -e, depending on the verb class) while third person plural agreement surfaces as a disyllabic suffix (-ano, -ono, depending on the verb class). As a result, any item forced toward a plural NP would have been one syllable (and two characters) heavier than an item forced toward a singular NP. While one syllable extra is much smaller than the weight manipulation used in other studies (e.g., Fernández 2003) and could have been controlled through counterbalancing, it could have potentially introduced noise into our data and would have necessitated a complication of our experimental design. Gender marking on a secondary predicate allowed us to avoid this issue as the gender marking on the secondary predicate is always monosyllabic (-o, -a).

To affect PR availability within the PR/RC items (β), we also manipulated the matrix predicate. In the PR/RC conditions, the matrix predicate was always a verb of perception (e.g., *vedere* - 'to see') or pseudo-perception (e.g., *filmare* - 'to film') which may take PRs. In the non-perceptual conditions, the matrix predicate was always a non-perception verb (e.g., *vivere con* - 'to live with') which does not admit PRs. This allowed us to maintain comparability with the sentence interpretation task reported in chapter 3.

Both the attachment and predicate manipulations were *within* items. This meant that for each item, 4 sentences were constructed, one per condition. Aside from the matrix predicate and the gender marking on the embedded secondary predicate, the four sentences for each item were identical. An example of a potential item is given in (66).

(66) a. *Gianni* ***vive con*** *il cognato$_M$ della signora$_F$ che ballava ubriaco**$_M$** nella fontana.*

b. *Gianni* ***vive con*** *il cognato$_M$ della signora$_F$ che ballava ubriaca**$_F$** nella fontana.*

c. *Gianni* ***ha visto*** *il cognato$_M$ della signora$_F$ che ballava ubriaco**$_M$** nella fontana.*

d. *Gianni* ***ha visto*** *il cognato$_M$ della signora$_F$ che ballava ubriaca**$_F$** nella fontana.*

'Gianni (saw/lives with) the brother-in-law of the woman (that was) dancing drunk in the fountain.'

4.2.2 Pronominal items

All pronominal items were of the template sketched in (67).

(67) NP_1 - Matrix Pred. - NP_2 - *mentre* ('while') - Pron. - Embedded Pred. - Secondary Pred. - ...

Within these items, we manipulated the pronominal form of the embedded subject (null or overt). To force pronominal reference even in the case of null subjects, we used the same gender manipulation as in

the PR/RC items (α, δ). Namely, we included a secondary predicate within the embedded CP. Given NP_1 and NP_2 always differed in gender, this was then used to force [±TS] readings. As we did not need complex NPs in the pronominal items, all NPs were simple proper names (e.g., Valeria, Adriano). On the one hand, this means the gender manipulations between our two stimuli sets were slightly different (recall that in the PR/RC items, we avoided any NPs which were lexically specified for gender). On the other hand, the use of proper names (i) allowed us to bypass any potential semantic biasing due to the lexical content of the NPs and (ii) are more felicitous in our decontextualised sentences. As with the PR/RC items, the gender of NP_1 was counterbalanced *between* items while the topic and pronoun manipulations were *within* item. Thus again for each item, we constructed 4 sentences. An example of a pronominal item is given in (68).

(68) a. *Valeria$_F$ ha salutato Adriano$_M$ mentre* ***lei*** *tornava affamata**a**$_F$ dalla palestra.* [Italian]

b. *Valeria$_F$ ha salutato Adriano$_M$ mentre* ***lui*** *tornava affamat**o**$_M$ dalla palestra.*

c. *Valeria$_F$ ha salutato Adriano$_M$ mentre* ***ø*** *tornava affamata$_F$ dalla palestra.*

d. *Valeria$_F$ ha salutato Adriano$_M$ mentre* ***ø*** *tornava affamat**o**$_M$ dalla palestra.*

'Valeria greeted Adriano while (she/he) was coming back hungry from the gym.'

As a final note on the pronominal items, it should be pointed out that our gender manipulation is not fully uniform with the PR/RC items. This is because overt singular pronouns in Italian manifest gender. As such in (68a, 68b) sentences are already disambiguated by the overt pronoun. Conversely, in (68c, 68d, and 66), the sentences are still ambiguous until encountering the secondary predicate. Although this is a disadvantage of using secondary predicates to disambiguate our items, items disambiguated by number would face a related issue. That is to say, we would independently need to model the reading times of null and overt pronominal items separately as in Chamorro et al. (2015a). This is because in items with overt pronouns (*lui, lei* - 'he, she') the window/region containing the embedded predicate will always be three characters longer than in items with null pronouns.

4.2.3 Segmentation

For use in the self-paced reading and eye-tracking-while-reading tasks, items were segmented as indicated in (69, 70). These 7 windows were used to define the windows in the self-paced reading task and the interest areas in the eye-tracking-while reading task. In PR/RC items, region 1 consisted of NP_0. Region 2 introduced the matrix predicate. Regions 3 and 4 consisted of NP_1 and NP_2. Region 5 included the complementiser and embedded predicate. Region 6 - the critical region - consisted exclusively of the disambiguating secondary predicate. Finally, region 7 was a follow-up region. This was included to push the critical region off the edge. In the pronominal items, region 1 consisted of NP_1. Region 2 introduced the matrix predicate. Region 3 contained NP_2 and region 4 consisted of the adverbial *mentre* ('while'). Region 5 contained the pronoun (when overt) and embedded predicate. Regions 6 and 7 were then identical to the PR/RC items. Regardless of item type, articles and prepositions were always included in the same region as the following NP.

(69)

Gianni	*ha visto*	*il cognato$_M$*	*della signora$_F$*	*che ballava*	*ubriaco$_M$*	*nella fontana.*
Gianni	has seen	the brother-in-law	of.the woman	that danced	drunk	in.the fountain
1	2	3	4	5	6	7

'Gianni saw the brother-in-law of the woman (that was) dancing drunk in the fountain.' [Italian]

(70)

Valeria$_F$	*ha salutato*	*Adriano*$_M$	*mentre*	*lei tornava*	*affamata*$_F$	*dalla palestra.*
Velaria	has greeted	Adriano	while	she was returning	hungry	from.the gym
1	2	3	4	5	6	7

'Valeria greeted Adriano while she was coming back hungry from the gym.'

4.2.4 Interim summary

The preceding section discussed the construction of our experimental items guided by a series of *desiderata*. By manipulating the matrix predicate within PR/RC items, we are able to control the syntactic availability of PRs (β). By exploiting gender marking on a secondary predicate, we are also able to force attachment/co-reference within both the PR/RC and pronominal items in a maximally uniform way (α, δ). Moreover, by using our gender manipulation rather than number marking on the tensed element in the embedded CP, we avoid the issue of phonological weight (γ). One issue which could not be explicitly controlled during the construction of the items, however, is potential semantic biases. As the semantic content of the embedded CP may affect attachment and indeed is sometimes used to disambiguate items (e.g., Dussias 2004), this should be accounted for. To that end, we now turn to our norming procedure.

4.3 Evaluation of the experimental items

As noted above (γ), we wanted to minimise potential influences from non-syntactic factors. In particular, we wanted to ensure that all the items were as natural and ambiguous (setting aside our gender disambiguation) as possible. Given we live in a gendered world, it is not inconceivable that some of our potential embedded predicates may have been biased to attach to a masculine or feminine NP (e.g. *che tornava affamat-o/-a dalla palestra* - 'coming back hungry from the gym'). As we cannot know *a priori* what gendered biases speakers may have, we conducted a norming study to elicit naturalness judgements. As some of our potential items may have been judged as unnatural or containing unwanted biases, we initially created 96 (48 per type) items. We then reduced the potential items based on the results of the norming study to 64 (32 per type; cf. Chamorro et al. 2015a, Aguilar et al. 2021) removing any items found to exhibit a gender bias or to be unduly unnatural.

4.3.1 Consultation with native linguists

As the author is not a native speaker of Italian, all critical items were first presented to two native Italian linguists (one for the PR/RC items and one for the pronominal items) who were made explicitly aware of the aims of the study. These informants were asked to check that: (i) all the sentences were grammatical (ii) all of the relevant readings were possible (e.g. for PR/RC items the informant was asked to judge whether the HA RC, LA RC, and HA PR readings were all accessible), (iii) all the items sounded natural, and (iv) that there were no obvious semantic biases in their interpretation. We opted to use linguists for this step to ensure that when judging the PR/RC items the informant could reliably discriminate between the HA PR and HA RC readings.

Responses from the two informants resulted in some minor lexical substitutions to improve the naturalness of the items. It also led to the exclusion of 3 pronominal items due to apparent semantic biases. As the naturalness and ambiguity judgements by the native informants were informal, the output from this step was then adapted for norming study with native Italian speakers.

4.3.2 Material for the norming study

The norming procedure for the present work was adapted from Hemforth, Fernandez, Clifton, Frazier, Konieczny & Walter (2015). For each potential item, we first isolated the embedded CP. We then formed 2 sentences from that lexical material, corresponding to the 2 potential subject readings for the original embedded CP.[54] In the case of the pronominal items such as (71a), this consisted of replacing the (null/overt) pronoun with the two potential referents from the matrix CP in the original item (71b, c). In the case of the PR/RC items such as (72a), this consisted of removing either just the complementiser (72b), or the complementiser, NP_1, and the preposition *di* ('of', 72c). In all items, the gender marking on the secondary predicate was made to agree with the now-subject NP. Sentences were distributed across two lists such that each participant saw only one sentence per pair, half in each condition. Participants never saw the original complex sentences (71a, 72a).

(71) Pronominal Norming Example

a. *Valeria ha salutato Adriano mentre (lei$_F$ / lui$_M$ / ø) tornava affamat-o$_M$/-a$_a$ dalla palestra.*
Valeria has greeted Adriano while she he *pro* came.back hungry from.the gym [Italian]

'Valeria greeted Adriano while she/he was coming back hungry from the gym.'

b. *Valeria tornava affamata dalla palestra.*

'Valeria was coming back hungry from the gym.'

c. *Adriano tornava affamato dalla palestra.*

'Adriano was coming back hungry from the gym.'

(72) PR/RC Norming Example

a. *Gianni ha visto il cognato$_M$ della signora$_F$ che ballava ubriac-o$_M$/-a$_F$ nella fontana.*
Gianni has seen the brother-in-law of.the woman that danced drunk in.the fountain

'Gianni saw the brother-in-law of the woman (that was) dancing drunk in the fountain.' [Italian]

b. *Il cognato della signora ballava ubriaco nella fontana.*

'The brother-in-law of the woman danced drunk in the fountain.'

c. *La signora ballava ubriaca nella fontana.*

'The woman danced drunk in the fountain.'

The decision to test only the plausibility of the 2 possible subject readings of the embedded CPs was motivated by the fact that the inclusion of the matrix predicate would have necessarily introduced syntactic biases. Violation of these biases (e.g. LA disambiguated perceptual items) has been previously observed to affect judgements. Looking at items temporarily ambiguous between PR- and RC-parses, Pozniak et al. (2019) reported that when their matrix predicate could introduce a PR, but the item was later disambiguated such that PR was no longer locally available, it resulted in a degradation of acceptability judgements. Convergent results have also been reported in Fernandes et al. (2018). Given that the point of the norming study is to

[54] As the norming study was conducted to identify items with potential semantic/plausibility biases in the interpretation of the embedded CPs, we temporarily collapse the HA RC and HA PR distinction. Regardless of the semantic type returned by the embedded CP, the basic event expressed remains unchanged.

control for any non-syntactic biases, the inclusion of the matrix predicate would muddle the results and be counterproductive. A similar argument can be made for the pronominal items. As noted above, speakers of consistent-null-subject languages exhibit clear interpretive biases for null and overt pronouns. When these biases are violated, speakers rate items as less acceptable (Chamorro et al. 2015a).

In addition to the PR/RC and pronominal items discussed above, the norming study also included nonsensical filler items (N = 20) such as (73). These were included as an attention control. That is to say, if participants were to rate highly implausible filler items as natural as the PR/RC or pronominal items, it would likely indicate that they did not understand the task, or were not paying sufficient attention. Conversely, if participants were to rate the fillers as significantly less natural than the PR/RC or pronominal items, this would indicate that they were attentive and sensitive to the plausibility of the items. As such, this would lend credence to the idea that we can use their judgements to evaluate the naturalness, and by extension the ambiguity, of the critical items.

(73) *Pietro coltivava le nuvole disteso sul letto.* [Italian]
Pietro grew the clouds lying on.the bed
'Pietro grew clouds lying in bed.'

4.3.3 Participants

A total of 30 participants (14 men, 16 women) were recruited. As with the experimental samples reported in the other chapters, all participants (i) were native Italian speakers, (ii) had lived in Italy from birth until at least the age of 16, (iii) had no diagnosed language-related disorders, and (iv) reported growing up monolingually.[55] At the time of testing, all participants were living in Italy. Due to an issue with data logging, one participant's data was lost. Of the remaining 29 participants, the average age was 43.34 years (SD = 7.72 years). Like the experimental groups reported in the other chapters, participants for the norming study were asked how many languages they spoke at or above an 'intermediate' level as well as for how many hours they used each of these languages in a typical day. This indicated that all of the remaining 29 participants spoke English to some degree. To quantify how much of each language a participant used, we calculated their percentages for each language. Table 4.1 presents group means and standard deviations for the use of Italian, English, and 'Other' as percentages of a typical day. Values were similar to the control groups for the sentence interpretation (chapter 3), self-paced reading, and eye-tracking studies (chapters 5 and 6).

	Mean	SD
Italian	81.00%	18.65%
English	12.53%	10.55%
Other	6.46%	16.06%

Table 4.1: Language use as percentages of a typical day for the participants of the norming study.

4.3.4 Procedure

Participants were recruited via Prolific Academic and received £3.75 for their participation. The experiment was carried out via PCIbex (Zehr & Schwarz 2018). After providing informed consent and completing the

[55] As with the other experimental samples, even though all participants reported growing up monolingually, when explicitly questioned about exposure to 'dialect' as a child (for example with one's grandparents) a portion of our sample indicated some exposure (N = 8).

language background questionnaire, participants then completed the naturalness judgement task. In this task, items were presented one at a time. Under each item, there was always the following question: 'From 1 (very unnatural) to 5 (very natural), how natural does the sentence sound to you?' An example of an item is presented in (74).

(74) *Il cognato della signora ballava ubriaco nella fontana.* [Italian]
the brother-in-law of-the woman danced drunk in-the fountain

'The brother-in-law of the woman was dancing drunk in the fountain.'

Da 1 (molto innaturale) a 5 (molto naturale), quanto ti suona naturale la
from 1 very unnatural) to 5 very natural how you-DAT sounds naturale the
frase?
sentence

'From 1 (very unnatural) to 5 (very natural), how natural does the sentence sound to you?'

Prior to starting the task, 'very natural' was explicitly defined to the participants as meaning that 'they might hear such a sentence in a normal conversation with other native speakers of Italian.' To respond, participants pressed the corresponding number key on their keyboard. This caused the sentence to disappear and the following item to take its place. Prior to data collection, ethical approval was obtained from the Ethics Committee of the Faculty of Modern and Medieval Languages and Linguistics at the University of Cambridge.

4.3.5 Results

Figure 4.1 presents the average rating for the three global item types. While participants rated the filler items as highly unnatural (mean = 1.27, SD = 0.86), the pronominal (mean = 4.70, SD = 0.72) and PR/RC items (mean = 4.27, SD = 1.03) were rated as quite natural.

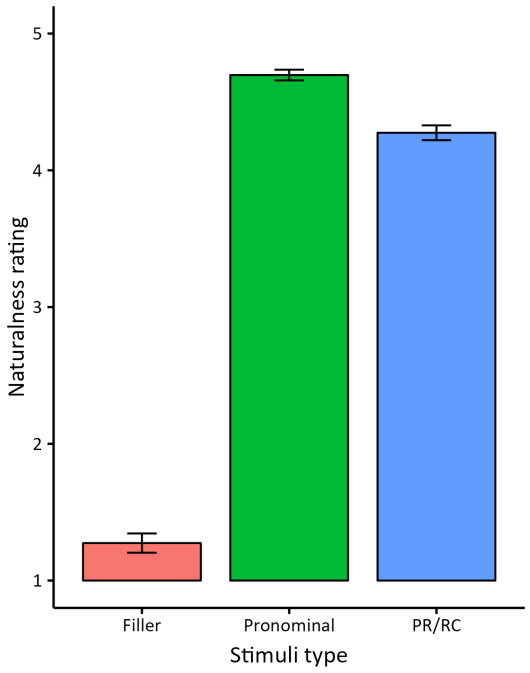

Figure 4.1: Global naturalness rating (on a scale from 1 'very unnatural' to '5 very natural') for norming items with 95% confidence intervals.

To see if participants were sensitive to the unnaturalness of the filler items, the three stimuli types were collapsed into fillers and non-fillers (i.e. PR/RC and pronominal items). These were then subjected to a mixed-effect regression with a fixed effect of *item type* (reference level: filler) and random effects of *item* and *participant*. The model (Table 4.2) indicated a significant effect of *item type* with the non-filler items being rated significantly higher than filler items.

	Estimate	Std. Error	*z*-value	*p*-value	
Intercept	1.27	0.11	12.01	<**0.001**	***
Item type	3.20	0.09	36.92	<**0.001**	***

Table 4.2: Model output for global norming ratings.

4.3.6 Discussion and selection of the final items

The fact the participants rated the non-filler items significantly higher than the filler items, suggests that they understood the task and were paying attention. As such, their rating of the PR/RC and pronominal items can be used to judge the global naturalness of the PR/RC and pronominal items as well as judge whether there are any potential semantic biases in any of the items.

As such, an unpaired t-test was conducted for each item. Any item for which there was a significant or marginal ($p < 0.1$) difference in the rating of the two sentence versions was deemed to be biased and as such removed. This resulted in the exclusion of 9 pronominal and 9 PR/RC items. The grand average rating for each item was then calculated (i.e. we collapsed the gender distinction), and the lowest rated items were removed until 32 PR/RC and 32 pronominal items remained such that the gender of NP_1 was balanced in both sets of items. The naturalness ratings for the embedded events in final PR/RC and pronominal items are presented in Table 4.3. The 4 original sentences for the final 64 items (i.e., the bi-clausal sentences in 66, 68) were then distributed across 4 lists such that each item appeared once and there were 8 items per condition. A list of the final critical items is reported in Appendix C.

	Mean	SD
PR/RC	4.32	0.99
Pronominal	4.81	0.54

Table 4.3: Naturalness ratings (on a scale from 1 'very unnatural' to 5 'very natural') for the final 32 items per type.

4.4 Chapter summary

This chapter motivated the need for the creation of a new set of experimental stimuli. It then sketched a series of *desiderata* and detailed how items were constructed to satisfy them. This chapter also described how these items were normed. As part of that process, we purposely created more items than necessary. These were then checked by native linguist consultants and presented to a group of native Italian speakers as part of a norming study. Based on the naturalness judgements provided by those participants we then reduced the original set of items to contain only those items that were maximally unbiased by plausibility and judged as most natural. As such, we can be more confident that any effects in the experiments described in the following chapters are due to syntactic parsing biases, not artefacts of the items.

5 Attrition and the Position of Antecedent Hypothesis in locally ambiguous items

5.1 Introduction

This chapter presents three experiments to investigate online pronominal bases in Italians living both in Italy and abroad. In section 5.3, we first present the results from a self-paced reading experiment with native Italian speakers still residing in Italy, using our new locally-ambiguous pronominal items. This was done to establish a baseline against which we could compare potential attriters. Results indicated the expected non-SpecIP bias for overt pronouns, but no bias for null pronouns. In section 5.4, we then present a sister experiment which used an eye-tracking-while-reading task. As in the first experiment, participants for the second experiment consisted only of native Italian speakers still residing in Italy. Given that we were unable to find a SpecIP bias for null pronouns in the first experiment, the eye-tracking-while-reading experiment is presented as a validation of this finding, albeit with a different methodology. Results from that experiment were convergent with those in the first self-paced reading experiment and suggest that we can use those results as a baseline against which to compare our group of Italians living abroad. To that end, section 5.5 presents a third experiment in which we presented native Italian speakers living in an English-speaking country with the same self-paced reading task as in section 5.3. This was done to explore potential attrition effects. Results indicated that attrition affected overt pronouns. However, *contra* Chamorro et al. (2015a), this appeared as a stronger non-SpecIP bias for the native Italian speakers living abroad.

5.2 Research questions

In this chapter, we focus on the following research questions:

(RQ_4) Within native Italian speakers still living in Italy, can we observe both PAS effects in our intrasentential pronominal items?

i. The null pronoun prefers an antecedent which is in SpecIP.

ii. The overt pronoun prefers an antecedent which is not in SpecIP.

(RQ_5) If so, how are these biases affected by attrition?

5.3 Experiment 1: Self-paced reading

5.3.1 Method

5.3.1.1 Participants

For the self-paced reading experiment, participants consisted of 66[56] native Italian speakers (Female = 33, Male = 33) who (i) had lived in Italy from birth until at least the age of 16, (ii) were living in Italy at the time of testing, (iii) had no diagnosed language related disorder, and (iv) reported growing up monolingually.[57] At

[56] This group is noticeably larger than groups elsewhere in the thesis due to the request of an anonymous reviewer.

[57] Despite all participants reporting growing up monolingually at the pre-screening, when specifically asked about potential exposure to 'dialect' as a child (for example with one's grandparents) 15 participants reported some exposure.

	Mean	SD
Italian	82.83%	14.41%
English	12.86%	11.95%
Other	4.31%	7.00%

Table 5.1: Mean language use by participant as percentages of a typical day with standard deviations.

the time of testing, the average age of participants was 43.42 years (SD = 7.49 years).[58] As in the experiment with ambiguous items reported in chapter 3, participants were asked to report how many languages they could speak at or above an 'intermediate' level. They were also asked to indicate how many hours they used each of these languages in a typical day. Of the 66 participants, all but 2 indicated that they could speak another language (and all but 4 could speak English). To quantify participants' language use, we calculated their percentages for each language (i.e., hours of language x / total hours reported for any language). Table 5.1 reports the means and standard deviations for their use of Italian, English, and 'Other.'

5.3.1.2 Stimuli

The final 32 pronominal items from the norming study reported in chapter 4 were used as the critical items. To recapitulate, the pronominal stimuli all contained a temporal adverbial clause which followed the matrix clause and contained the relevant pronoun. Within these items, we manipulated two factors: the pronominal form (null or overt) and the antecedent (SpecIP or non-SpecIP). To force particular antecedents, we used gender features as the two possible antecedents in the matrix clause always differed in gender. In items with a null pronoun, the disambiguating gender marking (-o.M, -a.F) only appeared on a secondary predicate immediately after the embedded verb. In items with an overt pronoun, gender marking also appeared on the pronoun itself (lui.M, lei.F). The gender of the possible antecedent in SpecIP of the matrix clause was counterbalanced *across* items. An example of an item in the four conditions is presented in (75) with vertical bars to indicate segmentation.

(75) a. Null SpecIP [Italian]

Valeria | ha salutato | Adriano | mentre | ø tornava | affamata | dalla palestra.

b. Overt SpecIP

Valeria	*ha salutato*	*Adriano*	*mentre*	*lei tornava*	*affamata*	*dalla palestra.*
Valeria.F	has greeted	Adriano.M	while	she came.back	hungry.F	from.the gym

'Valeria greeted Adriano while she was coming back hungry from the gym.'

c. Null non-SpecIP

Valeria | ha salutato | Adriano | mentre | ø tornava | affamato | dalla palestra.

d. Overt non-SpecIP

Valeria	*ha salutato*	*Adriano*	*mentre*	*lui tornava*	*affamato*	*dalla palestra.*
Valeria.F	has greeted	Adriano.M	while	he came.back	hungry.M	from.the gym

'Valeria greeted Adriano while he was coming back hungry from the gym.'

[58] As mentioned in footnote 42, although our participants are somewhat older than those in Chamorro et al. (2015a) or Martín-Villena (2023), they are more comparable to those in Tsimpli et al. (2004: Ianthi Tsimpli, p.c.). Moreover, in each experiment, our experimental and control groups are comparable in age. As such, age is highly unlikely to drive any potential group-level effects unlike what was reported in Kaltsa et al. (2015) who reported a 30 year age difference between groups.

Sentences for the pronominal items were distributed across 4 lists such that participants saw each item only once, with 8 sentences per condition. The experiment also included the 32 PR/RC items reported in chapter 6 as well as 40 unambiguous fillers for a total of 104 sentences.

5.3.1.3 Procedure

Prior to testing, we obtained ethics approval for the study from the Ethics Committee of the Faculty of Modern and Medieval Languages and Linguistics at the University of Cambridge. Participants were then recruited through Prolific Academic and paid £7.50 for their time. The experiment was run via PCIbex (Zehr & Schwarz 2018). After providing informed consent, participants filled out a language history questionnaire. This was followed by the self-paced reading task. For that task, participants were told that they would read a series of sentences in Italian piece by piece and then respond to a related question after each. Items were always presented alone on the participant's screen. They initially appeared as a series of underscores corresponding to the windows (boundaries indicated by '|'s in 75). In order to view the first window, participants pressed the space bar. When they had read this section, they pressed the space bar again. This caused the current window to disappear and the next window to become visible. After reading the final window, participants were presented with a polar comprehension question. For the pronominal and PR/RC items, this always asked about the subject of the non-matrix CP. In half of the items, the question asked about the first possible antecedent/attachment site and in half of the items, it asked about the second possible NP (counterbalanced *within* items). In the biclausal filler items, the comprehension questions asked about each of the CP in half of items. To respond, participants pressed 'F' or 'J.' This caused the experiment to pass to the next item. An example of a comprehension question for the item in (75) is presented in (76). To familiarise themselves with the procedure, participants were presented with practice items (N = 7) prior to testing.

(76) *Valeria tornava?* [Italian]

'Was Valeria coming back?'

F: Sì

'Yes'

J: No

'No'

We opted to include these comprehension questions as previous studies on different types of garden-path sentences have found that readers' initial/preferred parse may persist even after reanalysis/re-ranking (e.g., Christianson, Hollingworth, Halliwell & Ferreira 2001, Sturt 2007, Slattery, Sturt, Christianson, Yoshida & Ferreira 2013). This can be exploited to derive a secondary and indirect way of tapping parser biases, namely by comparing response accuracy after items in which we expect reanalysis and after items in which we do not. If accuracy is lower for the items that are expected to require reanalysis, this would support the idea that reanalysis has indeed taken place. This was particularly relevant for items with non-local disambiguations (e.g., the null pronominal items and the PR/RC items), but for consistency, we also report accuracy for the overt pronominal items.

5.3.1.4 Data cleaning and analysis

In the following analyses, we will consider the reading times of the disambiguating windows as well as the accuracy to the comprehension questions.

Due to a coding error, we lost one pronominal item across all lists (Pro 28, in Appendix C). We split the remaining 31 items with null and overt pronouns into two data sets to be analysed separately. This was

because the items were disambiguated at different points. Items with overt pronouns were disambiguated at the fifth window which contained the overt pronoun and embedded verb. Items with null pronouns were not disambiguated until the following region containing the secondary predicate (cf. 75a, 75b). To clean the reading times for the critical window of each set of items, we then coded as missing any item for which the analysed window, or any window prior to it, had an implausibly fast reading time (implemented as $<$ 200ms). Of the 31 analysed items, this affected 2.54% of the overt trials and 2.74% of the null trials. To identify potential outliers in the data we then calculated the median reading time and interquartile range (IQR) for each condition and coded as missing any value that lay 1.5 IQR above the third quartile. This affected a further 4.30% of the overt trials and 5.96% of null trials for a total data loss of 6.84% and 8.70% respectively. The remaining data was equally distributed across the 2 conditions (overt: $\chi^2(1, 953) = 0.03$; $p = 0.87$; null: $\chi^2(1, 934) = 0.15$; $p = 0.69$).

For the analysis of the comprehension question responses, we collapsed the two data sets back into one. This was because there were no differences in the comprehension questions, regardless of the difference in disambiguation points. For that analysis, we only considered those trials which were also included in the relevant reading time analyses.

To analyse the data we used the *lme4* package (Bates et al. 2015) in R (R Core Team 2022). We started by identifying the best random-effects structure. For this, we conducted a family of intercept-only models in which we considered the various random-effects structures of our theoretically relevant fixed predictors (i.e., random intercepts by *item* and *participant* as well as random slopes by *antecedent* and *pronoun* where relevant). From these, we selected the best fitting random effects structure (Matuschek et al. 2017) using the AIC. We then fit a base model with only the fixed effects that we were theoretically interested in using sum coding (i.e., -0.5, 0.5). For the modelling of reading times, this included only *antecedent* (negative level: SpecIP). For the modelling of the comprehension question responses we also included a potential interaction with *pronoun* (negative level: null). For consistency with the PR/RC data (collected in the same experimental session, and presented in chapter 6) we then conducted two further models which included *order* (centred over the experiment) as either a simple predictor or potential interaction term. If complicating the model reduced the AIC by 2 or more, the more complicated fixed-effect structure was maintained. For the analysis of the comprehension question responses, we then repeated this step with *response target* (negative level: yes). This was done to account for any potential effects of the acquiescence bias (Holbrook 2008) given the polar nature of our comprehension questions.

5.3.1.5 Hypotheses and predictions

The following hypotheses were drawn:

($H_{5.1}$) In items containing a null pronoun, participants will exhibit a SpecIP bias.

($H_{5.2}$) In items containing an overt pronoun, participants will exhibit a non-SpecIP bias.

For reading times then, we predict a simple effect of *antecedent* in both of the pronominal data sets. For items with null pronouns, we expect faster reading times when the pronoun co-refers with the DP in the matrix SpecIP than when it co-refers with a lower DP. For overt items, the opposite pattern is expected. However, we should stress that we do not predict that this will lead to an interaction of *antecedent* and *pronoun*. Recall that our design necessitates modelling the reading times for the two pronominal forms separately, due to the difference in the disambiguation points.

For the responses to the comprehension questions, assuming we can detect the persistence of initial/preferred parses, we would predict a similar pattern of results. Namely, we would predict higher accuracy for SpecIP items than non-SpecIP items when the pronoun is null. For overt items, the opposite pattern is expected.

Furthermore, as the response accuracy for both types of pronouns can be modelled together, we additionally expect an interaction of *antecedent* and *pronoun*.

As for order effects, we did not expect any specific interactions. This was because that predictor was included primarily as previous research has suggested that the PR-firstness effect is particularly susceptible to adaptation in these kinds of experimental contexts. Nonetheless, it is reasonable to expect that there may be a simple effect of *order* with participants readings faster as the experiment progresses.

Finally, should the final model include *response target*, we would expect participants to be more accurate when the target response was *si* ('yes') given the acquiescence bias.

5.3.2 Results

5.3.2.1 Reading times in window 5: Overt pronouns

Figure 5.1 graphically presents the mean reading times for items containing an overt pronoun in the fifth window, the window with the overt pronoun and embedded tensed element. On average items forced to co-refer with an antecedent in SpecIP were read slower (mean = 741.10 ms) than items forced to co-refer with an antecedent not in SpecIP (mean = 690.43 ms).

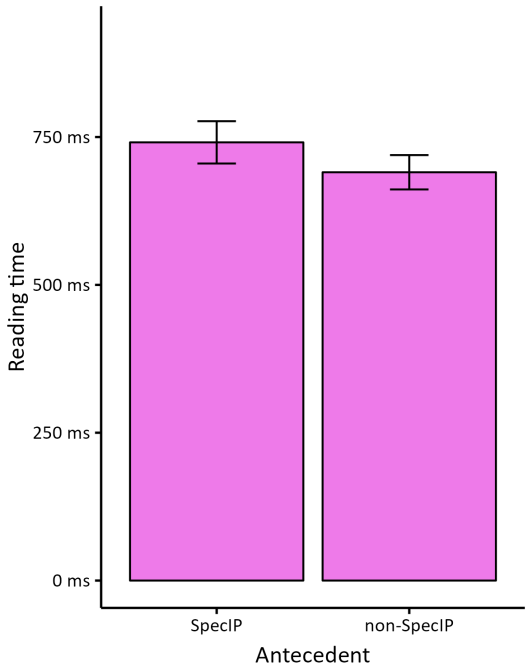

Figure 5.1: Global average reading times for the fifth window of items with overt pronouns with 95% confidence intervals.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.46	0.04	144.29	$< \mathbf{0.001}$	***
Antecedent	-0.05	0.02	-2.07	**0.04**	*
Order	-0.003	0.0004	-6.92	$< \mathbf{0.001}$	***

Table 5.2: Model output for the reading times of items with overt pronouns in the fifth window.

Reading times were logged and subjected to a family mixed-effects regression as described above. Model comparison indicated that the best fitting model included *order* as a simple effect (improvement to AIC >

2) but not an interaction term. The output for that model is reported in Table 5.2. That model indicated a significant effect of *antecedent* ($\hat{\beta}$ = -0.05; t = -2.07; p = 0.04) indicating that items forced to co-refer with a non-SpecIP antecedent were read more quickly. The model also indicated a significant effect of *order* ($\hat{\beta}$ = -0.003; t = -6.92; $p < 0.001$) with participants reading more quickly as the experiment progressed.

5.3.2.2 Reading times in window 6: Null pronouns

Figure 5.2 presents the average reading times for items containing a null pronoun in the sixth window, the window with the secondary depictive predicate. Average reading times were similar across the two conditions (SpecIP: 642.44 ms; non-SpecIP: 627.00 ms).

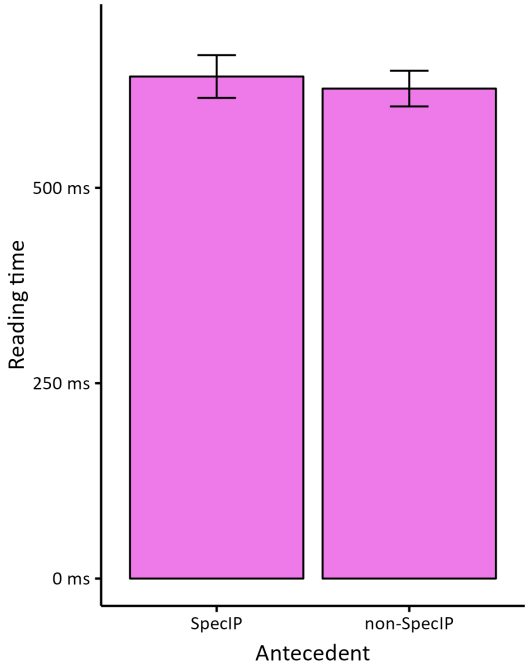

Figure 5.2: Global average reading times for the sixth window of items with null pronouns with 95% confidence intervals.

As with the items with overt pronouns, reading times were logged and subject to a family of mixed-effect regressions. Model comparison indicated that the best fitting model included an interaction between *antecedent* and *order* (improvement to AIC > 2). The output for that model is reported in Table 5.3. That model indicated a significant effect of *order* ($\hat{\beta}$ = -0.002; t = -5.94; $p < 0.001$) with participants reading more quickly as the experiment progressed. The effect of *antecedent* was not significant, but the interaction between the two was ($\hat{\beta}$ = 0.001; t = 2.10; p =0.04) indicating that the effect of *order* was smaller for non-SpecIP items. This is visually represented in Figure 5.3.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.39	0.04	163.44	$<$ **0.001**	***
Antecedent	-0.004	0.02	-0.21	0.83	
Order	-0.002	0.0003	-5.94	$<$ **0.001**	***
Antecedent:Order	0.001	0.0006	2.10	**0.04**	*

Table 5.3: Model output for the reading times of items with null pronouns in the sixth window.

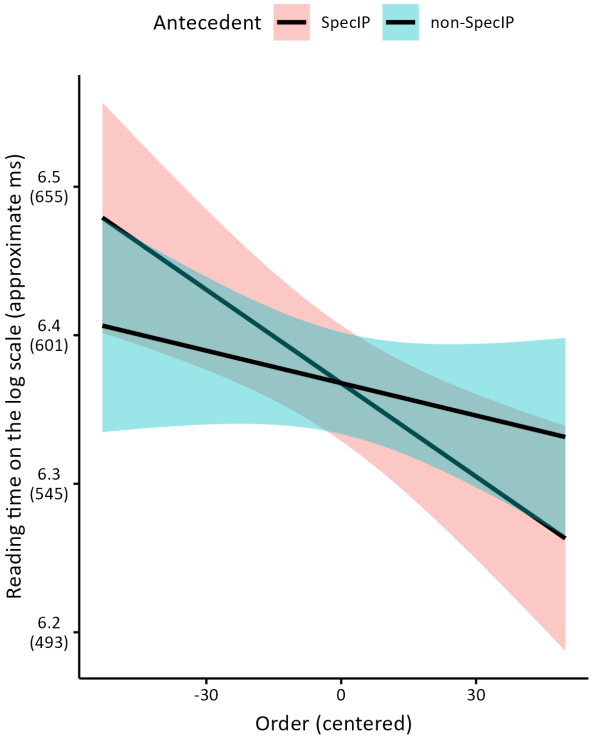

Figure 5.3: Average reading times over trial order on the log millisecond scale for the sixth window of items with null pronouns with 95% confidence intervals. This is plotted as a simple linear regression over the raw data for expositional purposes only.

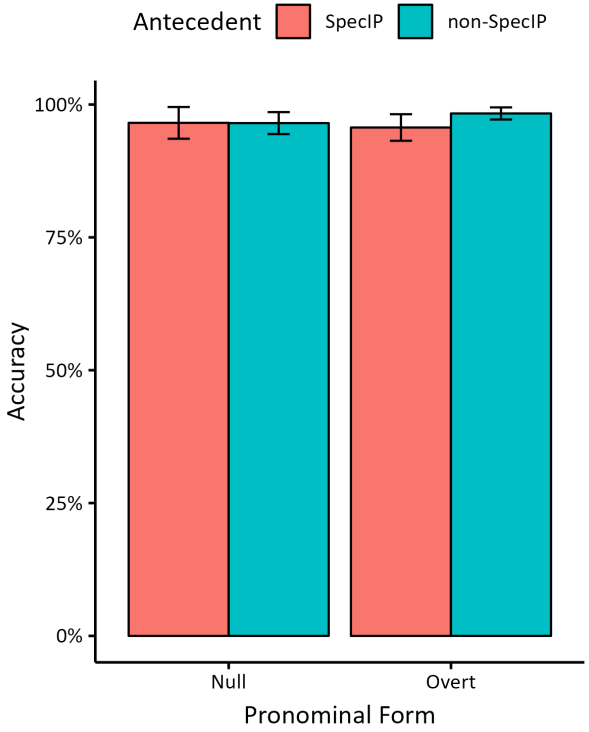

Figure 5.4: Average response accuracy (by participant) for null and overt pronominal items with 95% confidence intervals.

5.3.2.3 Response accuracy for comprehension questions: Null and overt pronouns

Figure 5.4 presents the average accuracy (by participant) for the comprehension questions after items with both null and overt pronouns. For items with null pronouns, participants gave the target answer 96.55% of the time when the pronoun co-referred with the possible antecedent in SpecIP and 96.50% of the time when it co-referred with the possible antecedent that was not. For overt pronouns, responses to SpecIP disambiguated items were 95.67% correct and responses to non-SpecIP disambiguated items were 98.32% correct.

Responses (coded as $\pm$ correct) were entered into a family of models as described above and model comparison indicated that the inclusion of *order* and *response target* did not improve the model fit (improvements to AIC < 2). As such we present the base model in Table 5.4. That model did not indicate a significant effect of *antecedent* or *pronoun*, nor did it indicate a significant interaction between the two.

	Estimate	Std. Error	*z*-value	*p*-value	
Intercept	4.37	0.03	13.12	$<$ **0.001**	***
Pronoun	0.28	0.29	0.96	0.34	
Antecedent	0.22	0.62	0.35	0.72	
Pronoun:Antecedent	1.05	0.58	1.811	0.07	.

Table 5.4: Model output for response accuracy for null and overt pronominal items.

5.3.3 Interim discussion

Above we have presented a self-paced reading experiment to establish a baseline against which we could later compare the experimental group. Specifically, we were interested in whether we could detect PAS effects, that is to say, a SpecIP bias for null pronouns and a non-SpecIP bias for overt ones. Starting with items containing overt pronouns, reading times for the disambiguating window were significantly longer when the pronoun was forced to co-refer with a potential antecedent in SpecIP of the matrix clause. This was as predicted ($H_{5.2}$) and is in line with what has been previously been reported for Italian. Turning to items with null pronouns, results indicated no bias *contra* ($H_{5.1}$). This was somewhat unexpected given previous results from Italian which have found a clear SpecIP bias (e.g., Carminati 2002, Vogelzang et al. 2019) and raises the question of whether something about our items may have affected the bias for null pronouns. Turning to the accuracy to comprehension questions, results did not indicate any accuracy penalty when the PAS biases were violated for either overt or null pronouns. This may indicate that our comprehension questions are not appropriate for detecting initial/preferred parses stemming from the PAS. We delay trying to interpret these null results until after the presentation of the sister eye-tracking-while-reading experiment.

5.4 Experiment 2: Eye-tracking-while-reading

5.4.1 Method

5.4.1.1 Participants

For this experiment, we recruited a different set of 27 native Italian speakers (Female: 18, Male: 9). As with the self-paced reading experiment, all participants (i) had lived in Italy from birth until at least the age of 16, (ii) were living in Italy at the time of testing, (iii) had no diagnosed language related disorder, and (iv) did not

report speaking any other languages at home as a child apart from some *dialetto* ('dialect').[59] At the time of testing, the average age of the group was 41.37 years (SD = 13.34). Table 5.5 reports the group's use of Italian, English and 'Other' as percentages of a typical day. The values are comparable to those reported for the self-paced reading participants (Table 5.1).

	Mean	SD
Italian	79.35%	16.52%
English	12.67%	13.38%
Other	7.98%	11.30%

Table 5.5: Mean language use by participant as percentages of a typical day with standard deviations.

5.4.1.2 Stimuli

The stimuli for the eye-tracking-while-reading experiment were identical to those in the self-paced reading experiment. Furthermore, we defined interest areas to be identical to the windows of the self-paced reading experiment. Examples are provided in (77) with vertical bars to indicate the critical interest areas for null and overt pronominal items.

(77) a. Null SpecIP [Italian]

Valeria ha salutato Adriano mentre ø tornava | affamata | dalla palestra.

b. Overt SpecIP

Valeria ha salutato Adriano mentre | lei tornava | affamata dalla palestra.
Valeria.F has greeted Adriano.M while she came.back hungry.F from.the gym

'Valeria greeted Adriano while she was coming back hungry from the gym.'

c. Null non-SpecIP

Valeria ha salutato Adriano mentre ø tornava | affamato | dalla palestra.

d. Overt non-SpecIP

Valeria ha salutato Adriano mentre | lui tornava | affamato dalla palestra.
Valeria.F has greeted Adriano.M while he came.back hungry.M from.the gym

'Valeria greeted Adriano while he was coming back hungry from the gym.'

5.4.1.3 Procedure

Prior to testing, we obtained ethics approval from Ethics Committee of the Faculty of Modern and Medieval Languages and Linguistics at the University of Cambridge. Participants were recruited via word of mouth and received €10 for their time. After providing informed consent, they proceeded to the eye-tracking-while-reading task. For this, we used an EyeLink 1000 Plus (SR Research). Stimuli were presented at a distance of 87 cm on a 48 by 27 cm screen with a resolution of 1920 by 1080 pixels. We recorded the movements of participants' right eyes with a 35mm lens at a sampling rate of 1000 Hz. Items always appeared as

[59] When specifically asked about potential exposure to 'dialect' as a child (for example with one's grandparents) seven reported some exposure.

a single (left-aligned) line of text in 32-point font (Calibri) and participants were tested individually in a sound-attenuated room in Italy.

Participants first read the instructions and then performed a 9-point calibration. They were then presented with a practice phase to familiarise themselves with the procedure. Before each item, they were required to fixate on a cross on the left-hand side of the screen in order to perform a drift check. If this were successful, the item would then appear immediately to the right of the fixation cross. If the drift check were unsuccessful, it would initiate a re-calibration sequence. After reading the item, participants pressed the space bar to view the comprehension question. As with the self-paced reading task, the comprehension question and possible answers (labelled 'F' and 'J') were presented alone on the screen. To respond, participants pressed the corresponding key.

5.4.1.4 Data cleaning and analysis

In the following, we consider eye-tracking measures for the relevant disambiguating regions as well as the accuracy to the comprehension questions.

To filter our eye-movement data, we first tried to merge any fixations under 80 ms with another fixation within 0.5 degree of visual angle. Where this was not possible, the fixations were removed. We also filtered out any fixation greater than 1000 ms. From the fixation data, we then extracted first pass, go past, and total time measures. The first pass (FP) is the sum of any fixations from the first time the region is entered from the left until the region is first exited (either to the left or the right). The go past (GP) is the sum of any fixations from the first time the region is entered from the left until it is exited to the right. This can include fixations to the left of the region. Finally, the total time (TT) is the sum of all fixations within a region regardless of whether it's the initial reading or a re-reading.

For the analysis of the disambiguating regions and the comprehension questions we first removed one pronominal item due to a coding error.[60] We then coded as missing any trials in which the critical region itself was not fixated. This affected 0.71% of the overt data and 0.96% of the null data. For the comprehension questions, we additionally excluded missing responses. This affected a further 0.36% of the recombined pronominal data (total exclusion = 1.19%).[61] For the FP and GP measures, we further coded as missing any trials for which the critical region was skipped in the initial reading (i.e., material to the right was read first). This affected a further 1.43% of the overt data (total loss = 2.14%) and 6.24% of the null data (total loss = 7.19%). Within each data set, the remaining data was equally distributed across the conditions (Overt reading: FP/GP: $\chi^2(1, 411) = 0.06$; $p = 0.81$; TT: $\chi^2(1, 417) = 0.002$; $p = 0.96$; Null reading: FP/GP: $\chi^2(1, 387) = 0.02$; $p = 0.88$; TT: $\chi^2(1, 413) = 0.002$; $p = 0.96$; Recombined pronominal accuracy: $\chi^2(1, 827) = 0.03$; $p = 1.00$).

To analyse the data, we again conducted families of models to (i) identify the best fitting random effects structure and (ii) account for potential effects that may obfuscate the pattern we are theoretically interested in. For the eye-movement data, we analysed the null and overt items separately as above. As such, the base models only contained *antecedent* as a fixed effect (sum coded, negative level: SpecIP). We then followed the model comparison procedure to explore potential effects of, and interactions with, *order* (centred over the experiment), selecting the best model using the AIC. For the pronominal comprehension data, we considered the two item types together. As such, the base model contained *antecedent* and *pronoun* (sum coded, negative level: null) as fixed predictors. The additional models for that data considered *order* and *response target* (sum coded, negative level: yes) in the same manner.

[60] Item 28, the same item removed from the self-paced reading experiment.
[61] The global loss due to non-fixations for the recombined null and overt pronouns was 0.84%

5.4.1.5 Hypotheses and predictions

Our hypotheses were the same as in the self-paced reading task.

($H_{5.1}$) In items containing a null pronoun, participants will exhibit a SpecIP bias.

($H_{5.2}$) In items containing an overt pronoun, participants will exhibit a non-SpecIP bias.

For eye-movement measures in the critical regions of the pronominal items, we expect an effect of *antecedent* for both pronominal forms. However, this is expected to surface in opposite directions. In items containing a null pronoun, we expect shorter durations for SpecIP disambiguations. In items containing an overt pronoun, we expect shorter durations for non-SpecIP disambiguations.

For the comprehension questions, we predicted greater accuracy when the disambiguations conformed to the PAS. Namely, we expected greater accuracy for null items when they were SpecIP disambiguated. Conversely, we expected greater accuracy for overt pronouns when they were disambiguated toward the non-SpecIP antecedent.

No specific predictions were made regarding *order* aside from a potential simple effect indicating that participants read more quickly as the experiment progressed. Finally, should the final model include *response target*, we would expect a simple effect with participants' responses being more accurate when the target response was *si* ('yes').

5.4.2 Results

5.4.2.1 Eye-movement measures in interest area 5: Overt pronouns

Table 5.6 presents the mean FP, GP, and TT for the fifth window which contained the overt pronoun and embedded tensed element. The same information is presented graphically in Figure 5.5. To analyse the data, we logged each measure separately and subjected them to families of mixed-effect regression as described above. Model comparison indicated that none of the best-fitting models contained *order* as a simple effect or potential interaction term. As such, the outputs for the base models are presented in Table 5.7.

	FP		GP		TT	
Antecedent	Mean	CI	Mean	CI	Mean	CI
SpecIP	443.05	31.79	571.01	58.66	947.03	87.91
Non-SpecIP	437.41	34.43	538.92	57.62	825.13	65.83

Table 5.6: Average FP, GP, and TT for overt pronominal items in ms for the fifth region with 95% confidence intervals.

The effect of *antecedent* was non-significant in the FP or GP models. In the TT model however, the effect was significant ($\hat{\beta}$ = -0.12; t = -2.61; $p < 0.01$) indicating that participants read non-SpecIP disambiguated items more quickly.

5.4.2.2 Eye-movement measures in interest area 6: Null pronouns

Table 5.8 presents the mean FP, GP, and TT for the sixth window which contained the disambiguating secondary depictive predicate. The same information is presented graphically in Figure 5.6. For the analysis, each measure was logged separately. The data was then subjected to families of mixed-effect regression as described above. Model comparison for the FP and GP measures indicated that the best-fitting models did not contain order as a simple term or potential interaction term. For the TT measure, the best fitting model

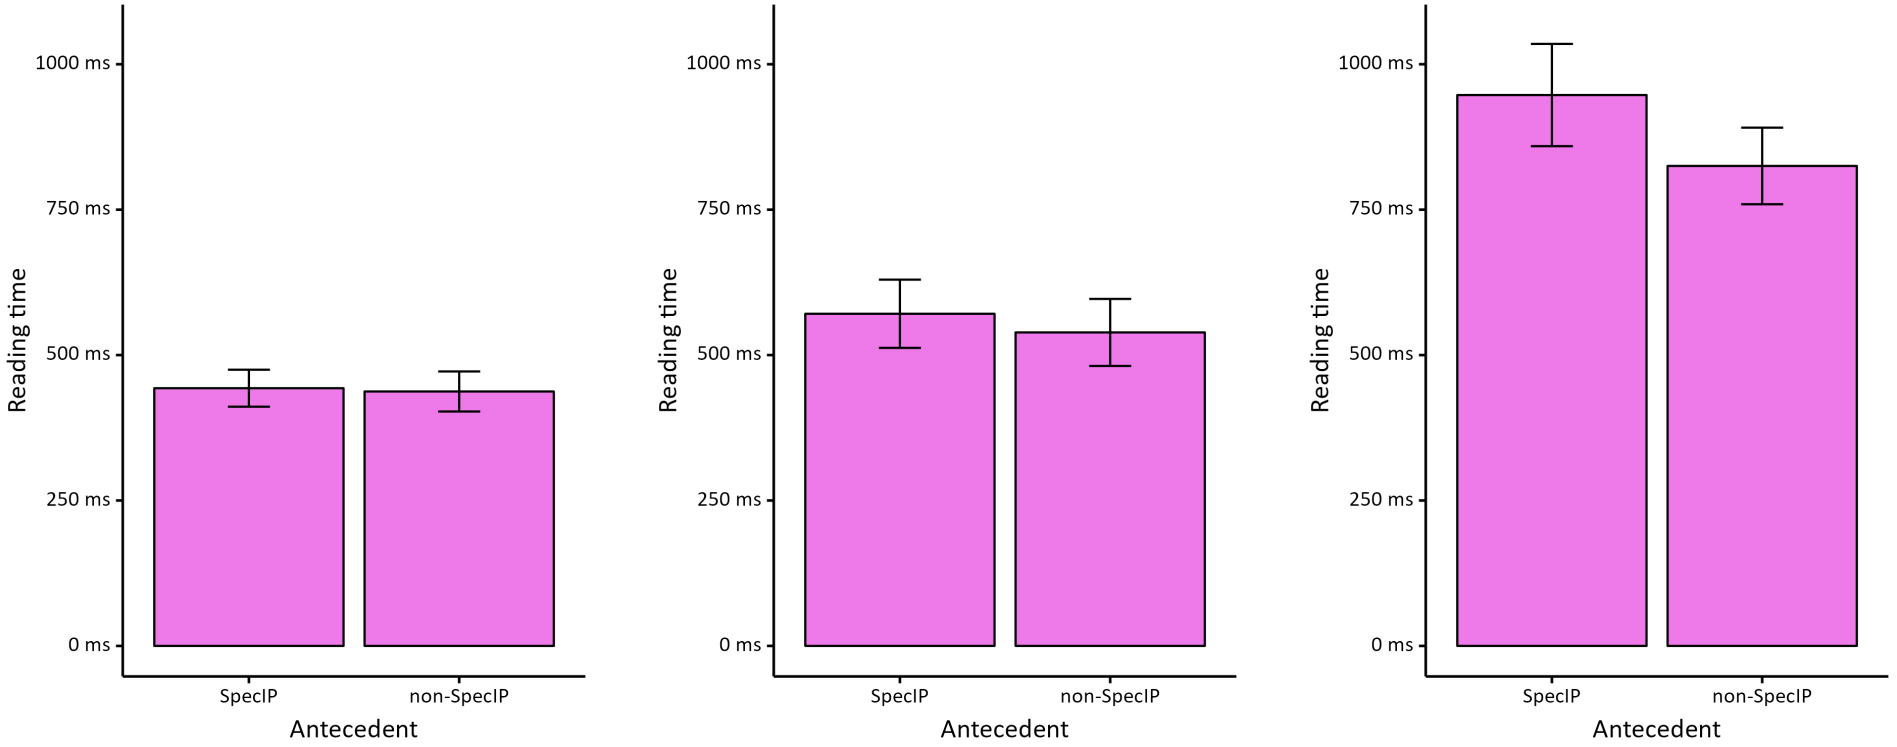

(a) FP (b) GP (c) TT

Figure 5.5: Eye-tracking measures for overt pronominal items in the fifth region with 95% confidence intervals.

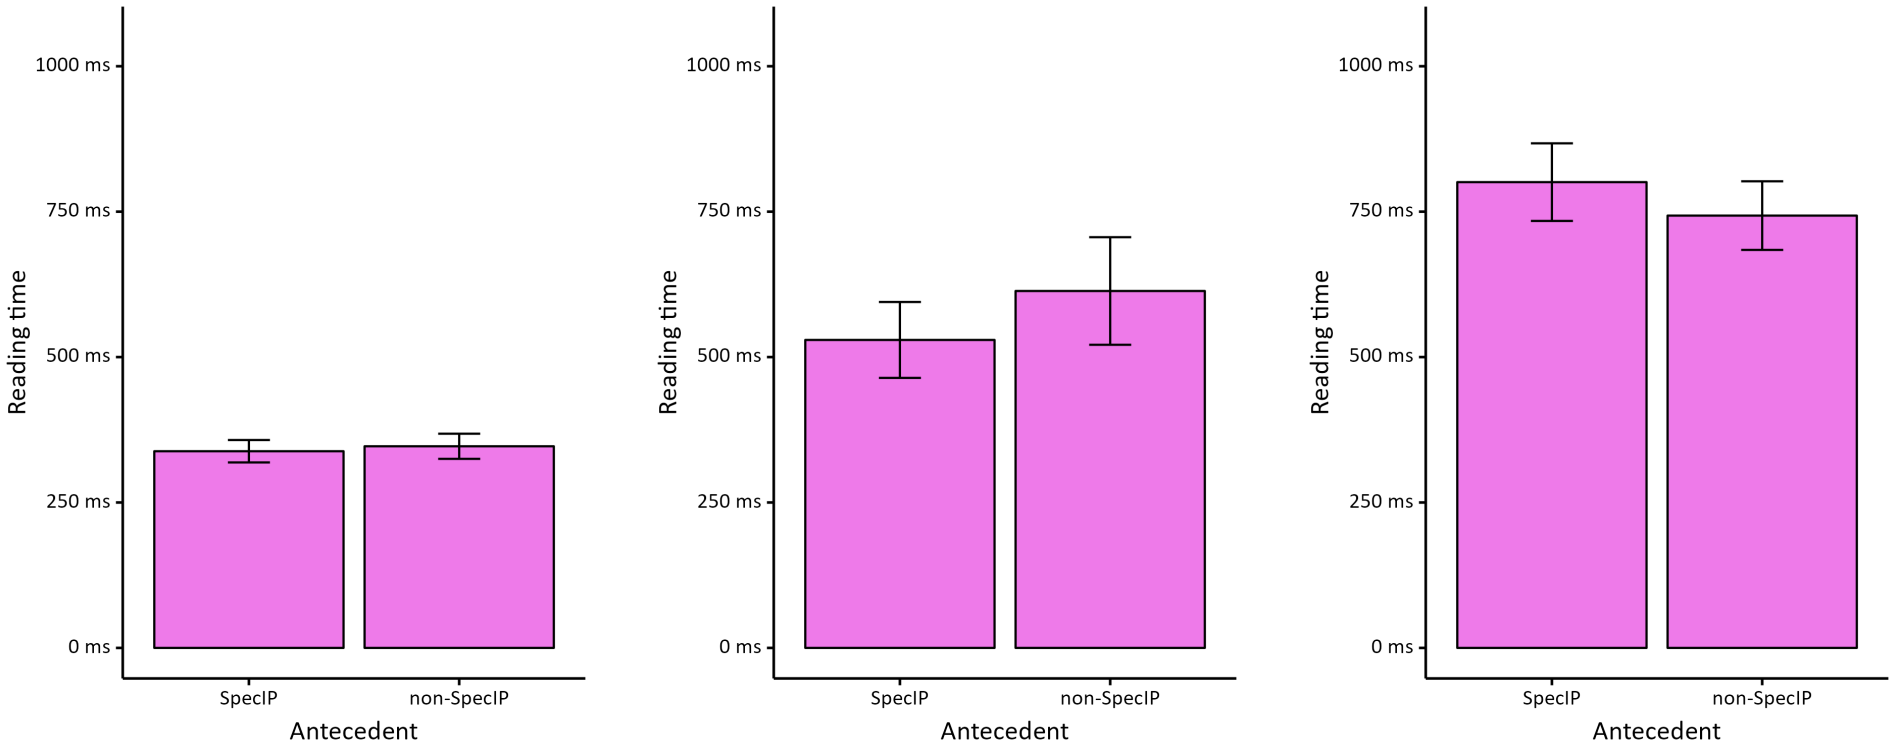

(a) FP (b) GP (c) TT

Figure 5.6: Eye-tracking measures for null pronominal items in the sixth region with 95% confidence intervals.

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	5.94	0.06	93.89	$<$ **0.001**	***
Antecedent	-0.04	0.04	-0.92	0.36	

(a) FP

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.11	0.06	94.33	$<$ **0.001**	***
Antecedent	-0.10	0.05	-1.89	0.06	·

(b) GP

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.59	0.09	73.01	$<$ **0.001**	***
Antecedent	-0.12	0.05	-2.61	$<$ **0.01**	**

(c) TT

Table 5.7: Model outputs for overt pronouns in the fifth region.

	FP		GP		TT	
Antecedent	Mean	CI	Mean	CI	Mean	CI
SpecIP	338.13	19.26	529.35	65.26	800.75	66.72
Non-SpecIP	346.64	21.61	613.61	92.47	743.22	58.98

Table 5.8: Average FP, GP, and TT for null pronominal items in ms for the sixth region with 95% confidence intervals.

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	5.74	0.04	131.22	$<$ **0.001**	***
Antecedent	0.01	0.04	0.32	0.75	

(a) FP

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.07	0.06	101.12	$<$ **0.001**	***
Antecedent	0.07	0.06	1.09	0.28	

(b) GP

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.47	0.07	87.41	$<$ **0.001**	***
Antecedent	-0.09	0.05	-1.80	0.07	·
Order	-0.003	0.0008	-3.73	$<$ **0.01**	**

(c) TT

Table 5.9: Model outputs for null pronouns in the sixth region.

included *order* as a simple effect (improvement to AIC > 2). The outputs for these best-fitting models are presented in Table 5.9.

The models indicated that the effect of *antecedent* was non-significant in the FP, GP or TT models. In the TT model however there was a significant effect of order ($\hat{\beta}$ = -0.003; t = -3.73; $p < 0.01$) with participants reading more quickly as the experiment progressed.

5.4.2.3 Response accuracy for comprehension questions: Null and overt pronouns

Figure 5.7 presents the mean accuracy to the comprehension questions for both null and overt pronominal items. For null pronominal items, participants selected the target answer for SpecIP disambiguated items 94.24%. When null pronouns were disambiguated toward the non-SpecIP antecedent accuracy was 98.48%. For overt pronouns, response accuracy for SpecIP disambiguated items was 97.53%. For non-SpecIP disambiguated items with overt pronouns, accuracy was 98.08%.

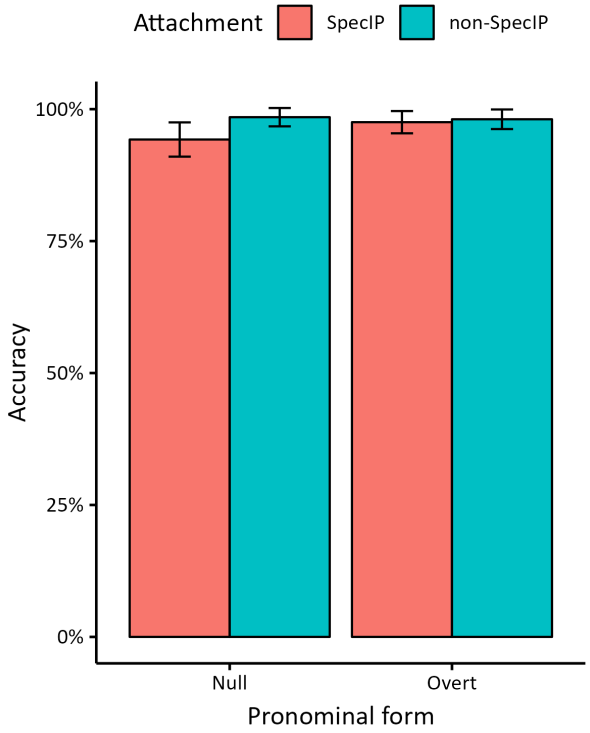

Figure 5.7: Average response accuracy (by participant) for null and overt pronominal items with 95% confidence intervals.

Responses (coded as $\pm$ correct) were entered into a family of logistic regression models as described above. Model comparison indicated that the best fitting model contained a simple effect of *order* (improvement to AIC > 2) but did not include *order* as an interaction term. Nor did it include *response target*. The best fitting model is presented in Table 5.10. The model only indicated a significant effect of *order* ($\hat{\beta} = 0.02$; $t = 2.71$; $p < 0.01$) with accuracy increasing over the course of the experiment.

5.4.3 Interim discussion

In this section, we have presented an eye-tracking-while-reading experiment to further investigate PAH biases in our new locally-ambiguous pronominal items. Specifically, we were interested in whether we could observe a SpecIP bias for null pronouns. Overall the results are very similar to those in the self-paced reading experiment above (section 5.3).

	Estimate	Std. Error	*z*-value	*p*-value	
Intercept	4.19	0.43	9.69	$<$ **0.001**	***
Antecedent	0.94	0.48	1.95	0.052	.
Pronoun	0.23	0.48	0.49	0.63	
Order	0.02	0.01	2.71	$<$ **0.01**	**
Antecedent:Pronoun	-1.24	0.96	-1.29	0.20	

Table 5.10: Model output for response accuracy for null and overt items.

Starting with eye-tracking measures for overt pronouns, we observed the expected non-SpecIP bias ($H_{5.2}$) in the TT measure, mirroring the bias found in the self-paced reading results.

Turning to eye-tracking measures for null pronouns, in no measure did we observe a significant bias for any antecedent. Although this is not the SpecIP bias one might expect from the PAH ($H_{5.1}$), this is in line with what was observed in the self-paced reading task. Moreover, some convergent results have also been reported by other studies working on null subject languages (e.g. Chamorro et al. 2015a, Kaltsa et al. 2015). Recall that Chamorro et al. (2015a) suggested that for their data, this may be related to the presentation order. Namely, Carminati (2002) originally studied the PAH in locally ambiguous items in which the embedded CP preceded the matrix CP. They chose that order specifically, arguing that the representation of the initial CP is likely to be kept in memory longer when it is not the matrix CP. As such, they assumed that upon encountering the pronoun, the parser will have greater access to structural information in the Embedded-Matrix order. This is relevant as the PAH is explicitly proposed as a structural bias. In both Chamorro et al.'s (2015a) study and the present study (which was modelled after Chamorro et al. 2015a for comparability of the experimental data in the next section), items instead appeared in the order Matrix-Embedded. However, such an explanation does not seem plausible for our Italian data. This is because other authors using different items in the same direction have observed clear evidence for a SpecIP bias with null pronouns (e.g., Vogelzang et al. 2019). As such, it is unclear why we were unable to detect a SpecIP bias for null pronouns. However, this appears to be a replicable result with our items.

As a final note on the two experiments with control participants, it appears that even though we can detect PAH effects in the reading times for overt pronouns, no such effect is observable within the accuracy to comprehension questions. As such, we interpret the null results for comprehension questions across the two experiments to indicate that such questions are not appropriate for detecting initial/preferred parses driven by the PAH (although as we will see in chapter 6 these questions are insightful for the PR/RC items). Regardless, the consistency between the two control experiments indicates that our original self-paced reading results, at least for the processing of overt pronouns, are a reliable baseline against which we can compare the experimental group in the next section.

5.5 Experiment 3: Self-paced reading again

5.5.1 Method

5.5.1.1 Participants

Experimental participants for the self-paced reading task consisted of 32 native speakers of Italian (Female = 24; Male = 8). As with the control group reported in section 5.3, all participants (i) had lived in Italy from birth until at least the age of 16, (ii) had no diagnosed language-related disorders, and (iii) reported growing

up monolingually.[62] Unlike the control group however, participants were living in a majority English-speaking country at the time of testing and had been there for a minimum of 5 years (mean = 12.14 years; SD = 7.65 years).[63] These participants were also prescreened to ensure that none had travelled back to Italy in the two months prior to testing to avoid any potential re-exposure effects (Chamorro et al. 2015a). At the time of testing the experimental group had a mean age of 39.44 years (SD = 7.73).[64] As with the control group, participants were asked how many languages they could speak at or above an 'intermediate' level. They were also asked to indicate for how many hours they used each of these languages in a typical day. To quantify participants' language use, we calculated their percentages for each language (i.e., hours of language x / total hours reported for any language). Table 5.11 reports the means and standard deviations for their use of Italian, English, and 'Other' alongside the values for the control group reported in section 5.3.

	Control		Experimental	
	Mean	SD	Mean	SD
Italian	82.83%	14.41%	16.16%	15.49%
English	12.86%	11.95%	82.58%	14.48%
Other	4.31%	7.00%	1.27%	3.67%

Table 5.11: Mean language use by participant as percentages of a typical day with standard deviations for the experimental and control groups.

As some previous work on attrition has restricted itself to 'near-natives' (e.g., Tsimpli et al. 2004), we had this group take the Cambridge Assessment General English quick placement test.[65] On average the group scored 22.16 / 25 (SD = 2.02) suggesting they are likely upper-intermediate to advanced L2 speakers of English. As such, the experimental group's prolonged residency in their L2 community, lack of recent re-exposure, frequent L2 use, attenuated use of the L1, and high L2 proficiency suggest this is a suitable group in which to look for attrition.

5.5.1.2 Stimuli

The stimuli were identical to those used for the control group. An example of the critical items (N = 32) in the null SpecIP condition is partially repeated in (78). As before, the experiment also included the 32 PR/RC items reported in chapter 6 as well as 40 unambiguous fillers for a total of 104 sentences.

(78) a. Null SpecIP [Italian]

Valeria | *ha salutato* | *Adriano* | *mentre* | *ø tornava* | *affamata* | *dalla palestra.*

b. Overt SpecIP

Valeria	*ha salutato*	*Adriano*	*mentre*	*lei tornava*	*affamata*	*dalla palestra.*
Valeria.F	has greeted	Adriano.M	while	she came.back	hungry.F	from.the gym

'Valeria greeted Adriano while she was coming back hungry from the gym.'

[62] Again, even though all participants reported growing up monolingually at the pre-screening, when specifically asked about potential exposure to 'dialect' as a child (e.g., with one's grandparents), 8 indicated some exposure.

[63] Note, as all participants reported growing up monolingually in Italy, the experimental participants qualify as first-generation immigrants (i.e., potential attriters), not second-generation immigrants (i.e., heritage speakers).

[64] As mentioned in footnote 42, although our participants are somewhat older than those in Chamorro et al. (2015a) or Martín-Villena (2023), they are more comparable to those in Tsimpli et al. (2004: Ianthi Tsimpli, p.c.). Moreover, in each experiment, our experimental and control groups are comparable in age. As such, age is highly unlikely to drive any potential group-level effects unlike what was reported in Kaltsa et al. (2015) who reported a 30 year age difference between groups.

[65] Items available at: https://www.cambridgeenglish.org/test-your-english/

5.5.1.3 Procedure

The procedure for the experimental group was identical to the control group except for the addition of the English placement test. Prior to testing, we obtained ethics approval for the study from the Ethics Committee of the Faculty of Modern and Medieval Languages and Linguistics at the University of Cambridge. Participants were then recruited through Prolific Academic and the experiment was run via PCIbex (Zehr & Schwarz 2018). After providing informed consent, participants filled out a language history questionnaire. This was followed by the self-paced reading task (window boundaries indicated by '|'s in 75, partial.). After each item, participants were presented with a polar comprehension question with the possible responses labelled 'J' and 'F.' After completing the self-paced reading task, participants were then presented with the English placement test.

5.5.1.4 Data cleaning and analysis

In the following analyses, we will consider the reading times of the disambiguating windows as well as the accuracy to the comprehension questions. For these analyses, we pool the data from the Italians still living in Italy reported in section 5.3 (henceforth the control group) and the new data from the Italians living in a majority English-speaking country (henceforth the experimental group).

Prior to analysis, we first cleaned the data from the experimental group. For that data, we removed one pronominal item across all lists (Pro 28, in Appendix C) due to the coding error noted above. We split the remaining 31 items for the experimental group into two data sets (null and overt) to be analysed separately due to the difference in the disambiguation point (cf. 75, partial a and b). To clean the reading times for the critical window of each set of items, we again coded as missing any item for which the analysed window, or any window prior to it, had an implausibly fast reading time (implemented as < 200ms). This affected 2.22% of the overt trials and 1.41% of the null trials. To identify potential outliers in the data we then calculated the median reading time and IQR for each condition and coded as missing any value that lay 1.5 IQR above the third quartile. This affected a further 2.82% of the overt trials and 5.24% of null trials for a total data loss of 5.04% and 6.65% respectively. The remaining data was equally distributed across the 2 conditions (overt: $\chi^2(1, 471) = 0.002$; $p = 0.96$; null: $\chi^2(1, 463) = 0.05$; $p = 0.82$).

For the analysis of the comprehension question responses, we collapsed the two data sets back into one and only considered the trials included in the relevant reading time analyses.

To analyse the data we used the *lme4* package (Bates et al. 2015) in R (R Core Team 2022). We followed the same procedure to identify the best random-effects structure. That is to say, we conducted families of intercept-only models in which we considered the various random-effects structures of our theoretically relevant fixed predictors. From these, we selected the best fitting random effects structure (Matuschek et al. 2017) using the AIC. We then fit a base model with only the fixed effects that we were theoretically interested in using sum coding (i.e., -0.5, 0.5). For the modelling of reading times, this included *antecedent* (negative level: SpecIP) and *group* (negative level: control). For the modelling of the comprehension question responses, we also included a potential interaction with *pronoun* (negative level: null). For consistency, we then conducted two further models which included *order* (centred over the experiment) as either a simple predictor or potential interaction term. If complicating the model reduced the AIC by 2 or more, the more complicated fixed-effect structure was maintained. For the analysis of the comprehension question responses, we then repeated this step with *response target* (negative level: yes) to account for any potential effects of the acquiescence bias (Holbrook 2008).

5.5.1.5 Hypotheses and predictions

The following hypotheses were drawn:

($H_{5.3}$) In items containing a null pronoun, overall participants will exhibit no bias.

($H_{5.4}$) In items containing an overt pronoun, the experimental group will exhibit a weaker non-SpecIP bias than the control group (cf., Chamorro et al. 2015a).

For reading times of items with null pronouns, given the results from our first two experiments, we expect no difference in reading times between when the pronoun co-refers with the DP in the matrix SpecIP and when it co-refers with a lower DP. Moreover, based on the results reported in Chamorro et al. (2015a), we do not expect an interaction between *antecedent* and *group*. Note however, if we do observe an interaction we might expect this to surface as a true SpecIP bias in the experimental group given the results reported in chapter 3.

For overt items, however, we expect a main effect of *antecedent* with faster reading times for non-SpecIP disambiguated items, which is modulated by an interaction with *group*. This prediction is empirically motivated by Chamorro et al. (2015a) who observed a weakening of the non-SpecIP bias for overt pronouns.

For the responses to the comprehension questions, no effects or interactions are expected given the results from the first two experiments (sections 5.3 and 5.4). Similarly, we did not make any specific predictions for order effects beyond a potential simple effect of *order* with participants reading faster as the experiment progresses.

5.5.2 Results

5.5.2.1 Reading times in window 5: Overt pronouns

Figure 5.8 presents the mean reading times for the fifth window of items containing an overt pronoun for the experimental group alongside the control group. On average the experimental group read items forced to co-refer with an antecedent in SpecIP slower (mean = 849.22 ms) than items forced to co-refer with an antecedent not in SpecIP (mean = 704.80 ms).

Reading times from both groups were logged and subjected to a family of mixed-effect regressions as described above. Model comparison indicated that the best fitting model contained *order* as a potential interaction term (improvement to AIC > 2). The output of that model is reported in Table 5.12.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.50	0.04	166.98	$<$ **0.001**	***
Antecedent	-0.09	0.02	-3.91	$<$ **0.001**	***
Group	0.08	0.07	1.00	0.32	
Order	-0.003	0.0003	-8.37	$<$ **0.001**	***
Antecedent:Group	-0.09	0.05	-1.92	0.06	·
Antecedent:Order	0.002	0.0007	2.65	$<$ **0.01**	**
Group:Order	-0.005	0.0007	-0.71	0.48	
Antecedent:Group:Order	0.003	0.001	2.00	**0.046**	*

Table 5.12: Model output for the reading times of items with overt pronouns in the fifth window for the experimental and control groups.

The model indicated a significant effect of *antecedent* ($\hat{\beta}$ = -0.09; t = -3.91; $p < 0.001$) with participants reading non-SpecIP disambiguated items more quickly. There was also a significant effect of *order* ($\hat{\beta}$ = -0.003; t = -8.37; $p < 0.001$) indicating that participants read more quickly as the experiment progressed. The effect of *group* was non-significant as were the interactions of *group* by *antecedent* and *group* by *order*. The model did however indicate a significant interaction of *antecedent* by *order* ($\hat{\beta}$ = 0.002; t = 2.65; p

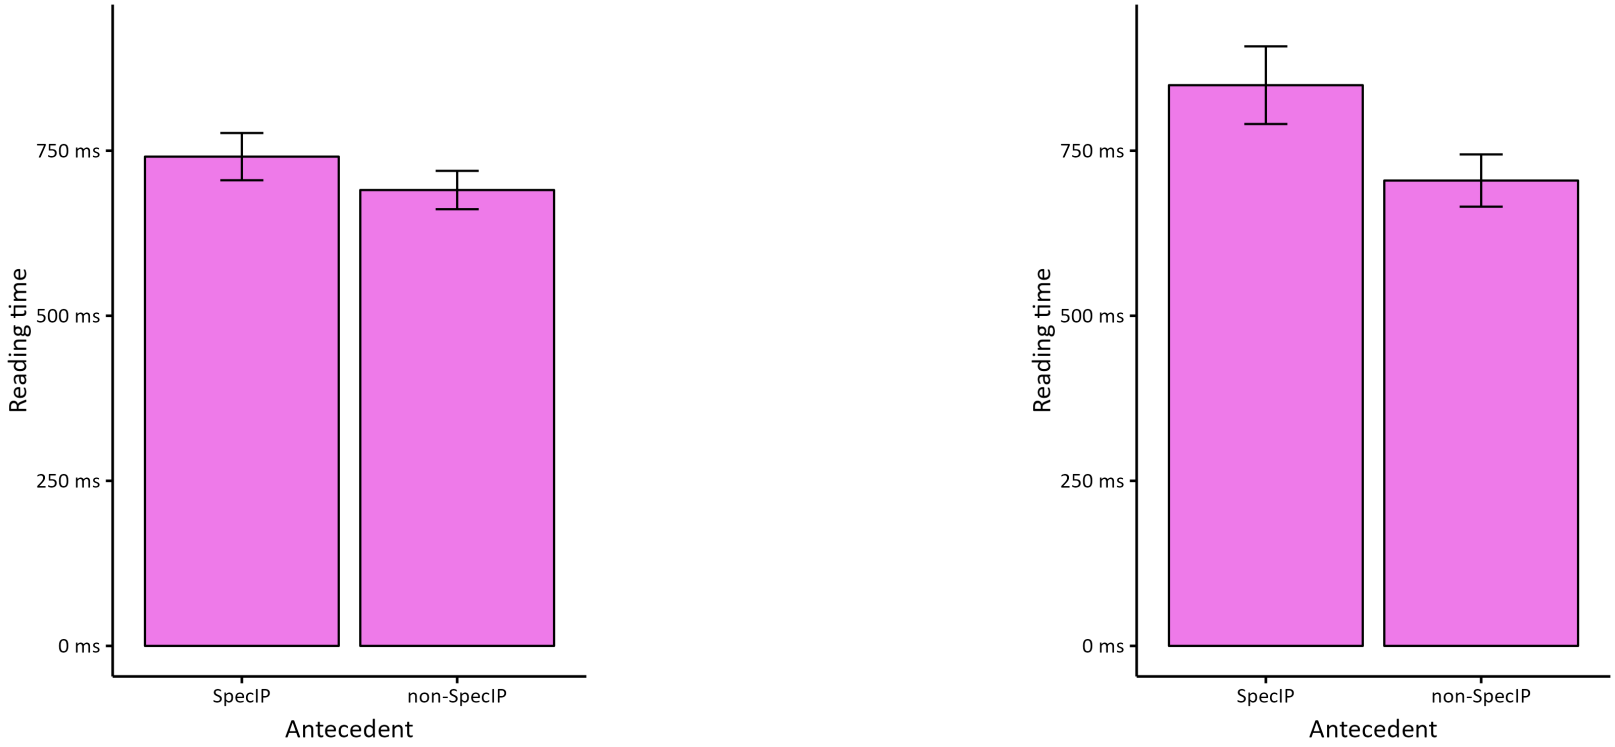

(a) Control

(b) Experimental

Figure 5.8: Global average reading times for the fifth window of items with overt pronouns for the control and experimental groups with 95% confidence intervals.

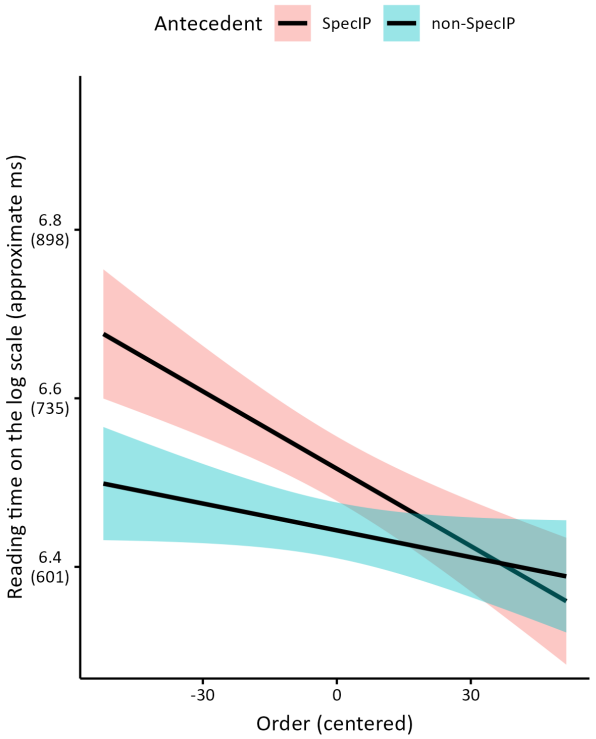

Figure 5.9: Average reading times by antecedent over trial order in log milliseconds for the fifth window of items with overt pronouns with 95% confidence intervals. This is plotted as a simple linear regression over the raw data for expositional purposes only. Note for legibility of the non-linear scale, the y-axis does not start at 0.

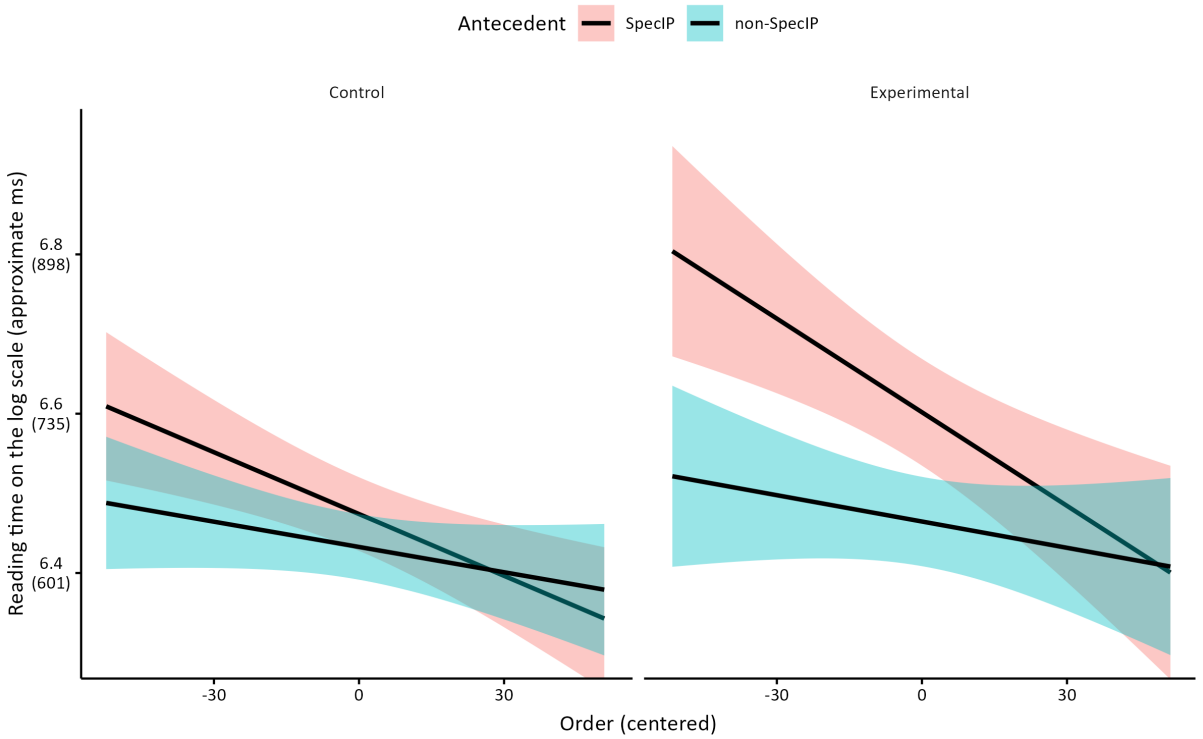

Figure 5.10: Average reading times by antecedent and group over trial order in log milliseconds for the fifth window of items with overt pronouns with 95% confidence intervals. This is plotted as a simple linear regression over the raw data for expositional purposes only. Note for legibility of the non-linear scale, the y-axis does not start at 0.

< 0.01). This indicated that the effect of *order* was smaller for non-SpecIP disambiguated items. That is to say, the reading time cost of violating the PAH bias for overt pronouns decreased over the course of the experiment. This is graphically represented in Figure 5.9. The model also indicated a significant interaction of *antecedent*, *group*, and *order* ($\hat{\beta} = 0.003$; $t = 2.00$; $p = 0.046$) indicating that the differential effect of *order* on the two possible disambiguations was larger in the experimental group.[66,67] This is graphically represented in Figure 5.10.

5.5.2.2 Reading times in window 6: Null pronouns

Figure 5.11 presents the mean reading times for the sixth window of items containing a null pronoun for the experimental group alongside the control group. On average the experimental group read items forced to co-refer with an antecedent in SpecIP similar (mean = 652.37 ms) to items forced to co-refer with an antecedent not in SpecIP (mean = 662.75 ms).

As above, reading times from both groups were logged and subjected to a family of mixed-effect regressions. Model comparison indicated that the best fitting model contained *order* as a simple predictor but not as a potential interaction term. The output of that model is reported in Table 5.13. The model indicated a significant effect of *order* ($\hat{\beta}$ = -0.002; t = -7.19; $p < 0.001$) with participants reading more quickly as the experiment progressed. The effects of *group* and *antecedent* as well as their interaction were non-significant.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.40	0.03	194.53	$<$ **0.001**	***
Antecedent	0.003	0.02	0.18	0.86	
Group	0.04	0.06	0.65	0.52	
Order	-0.002	0.0003	-7.19	$<$ **0.001**	***
Antecedent:Group	0.01	0.03	0.43	0.67	

Table 5.13: Model output for the reading times of items with null pronouns in the sixth window for the experimental and control groups.

5.5.2.3 Response accuracy for comprehension questions: Null and overt pronouns

Figure 5.12 presents the average accuracy (by participant) for the experimental group alongside the control group for reference. In both cases, accuracy is broken down by *pronoun* and *antecedent*. For items with null pronouns, experimental participants gave the target answer 94.68% of the time when the pronoun co-referred with the possible antecedent in SpecIP and 95.80% of the time when it co-referred with the possible antecedent that was not. For overt pronouns, responses to SpecIP disambiguated items were 96.88% correct and responses to non-SpecIP disambiguated items were 96.63% correct.

Responses (reference level: incorrect) were entered into a family of models as described above and model comparison indicated that the inclusion of *order* and *response target* did not improve the model fit (improvements to AIC < 2). As such we present the base model in Table 5.14. That model did not indicate

[66] In a simplified model that does not take *order* into consideration, the two-way interaction of *group* and *antecedent* is significant. The output for that model may be found in Table D.1.

[67] To follow up on the increased processing cost of violating the PAH within the experimental group (modulated by order) we conducted a series of exploratory analyses. For these, we only considered the data from the experimental participants and applied the same model selection procedure as with *order*. After identifying the best random effect structure, we then fit a base model with *antecedent* and *order* as fixed predictors. We then attempted to add *Length of Residence* (in the L2), *English exposure*, and *English proficiency* (as measured by the English placement test) as simple effects or potential interaction terms. Model comparisons indicated, that only the inclusion of *English proficiency* as a simple effect improved model fit. The output for that model may be found in Table D.2.

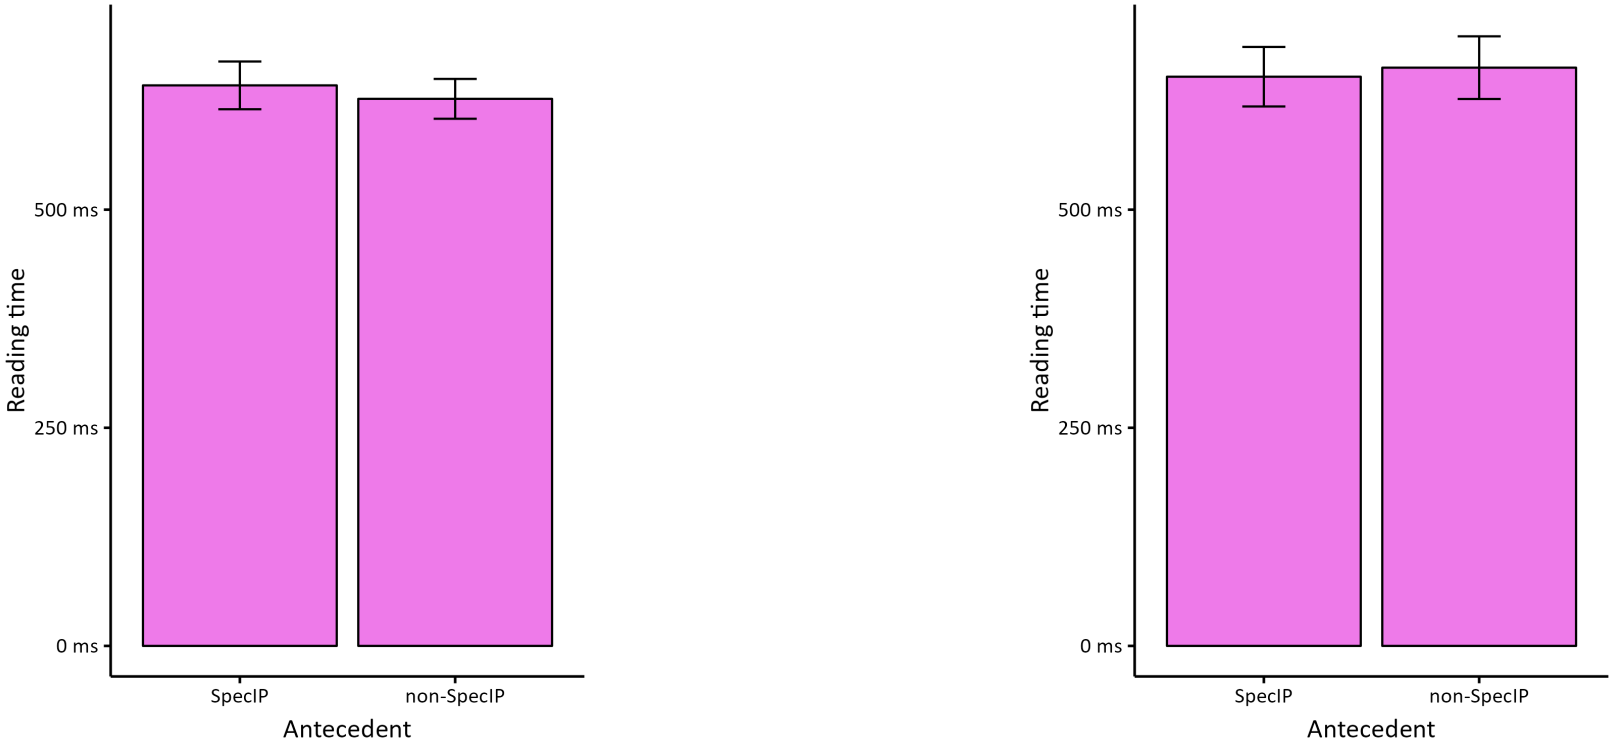

(a) Control

(b) Experimental

Figure 5.11: Global average reading times for the sixth window of items with null pronouns for the control and experimental groups with 95% confidence intervals.

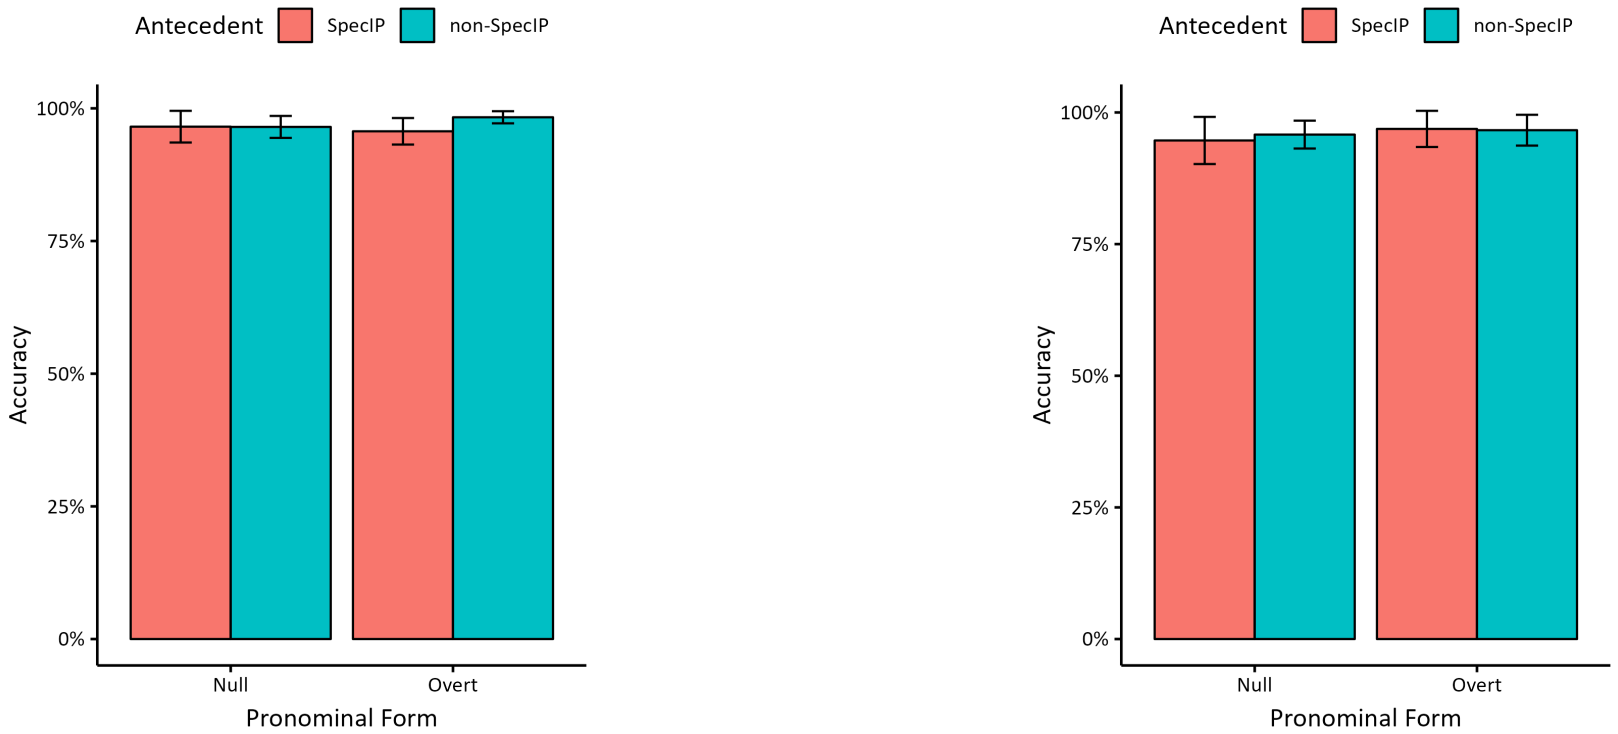

(a) Control

(b) Experimental

Figure 5.12: Average response accuracy (by participant) for the control and experimental groups broken down by pronominal type with 95% confidence intervals.

a significant effect of *antecedent*, *pronoun* or *group*, nor did it indicate any significant interactions between them.

	Estimate	Std. Error	*z*-value	*p*-value	
Intercept	4.26	0.28	15.45	$<$ **0.001**	***
Pronoun	0.38	0.23	1.62	0.11	
Antecedent	-0.16	0.52	-0.31	0.76	
Group	-0.003	0.37	-0.65	0.51	
Pronoun:Antecedent	0.32	0.70	2.65	0.48	
Pronoun:Group	0.19	0.46	0.41	0.68	
Antecedent:Group	-0.50	0.62	-0.80	0.42	
Pronoun:Antecedent:Group	-1.45	-1.57	2.00	0.12	

Table 5.14: Model output for response accuracy for the experimental and control groups.

5.6 Discussion

Above we have presented the experimental group (native Italian speakers living in a majority English-speaking country) for our self-paced reading task and compared their results against the control group from section 5.3. This was done to investigate how attrition affects pronominal processing with respect to the PAH.

Starting with the comprehension questions, the interpretation of the results is straightforward. Accuracy for the experimental group was at ceiling regardless of pronominal type or antecedent and their accuracy did not differ from the control group's. However, we have independently argued that our comprehension questions were not suitable for detecting preferred parses stemming from the PAH. Recall that the two experiments which focused on native speakers also found null results for the comprehension questions despite clear non-SpecIP biases for overt pronouns in online measures. Therefore, we take our questions to be unsuitable for detecting attrition, and as a result, we should refrain from interpreting the lack of a group effect.

Turning to the reading times for the overt pronominal items, results indicated that overall participants exhibited a non-SpecIP bias, with the caveat that this was affected by trial order. Given that the interaction of antecedent by order surfaced as a reduced effect of antecedent over the course of the experiment, we interpret this as adaptation to the experimental context. As such, we take reading times for items closer to the start of the experiment (where participants exhibit a clear non-SpecIP bias) to be more representative of how participants would process such structures in non-experimental contexts than reading times for items closer to the end of the experiment (where participants exhibit no bias, Figure 5.9). As for why this adaptation effect arose, we might first wonder whether there were insufficient non-critical items. This seems unlikely given that the relevant items (i.e., those with an overt pronoun) only accounted for 15.38% of items (16 / 104). Moreover, given our counterbalancing of the gender manipulation, we cannot attribute the adaptation effect to a 'highly reliable cue.' That is to say, upon identifying the genders of NP_1 and NP_2, participants still had no way of predicting which antecedent would be forced, unlike the tense manipulation in Pozniak et al. (2019) in which present tense was always incompatible with PR parses.

A salient alternative possibility is that the adaptation effect might relate to the nature of our experimental design itself. Within the same experiment session, we also included null pronominal and PR/RC items. All of those were locally ambiguous for a total of 61.54% of our items (64 / 104). At the beginning of

the experiment, participants could not have been aware of this commonality throughout our items. As the experiment progressed, however, participants were repeatedly garden pathed. Given the unreliability of their parsing biases in the specific experimental context, participants may have instead begun to rely more on other cues, in particular the disambiguating gender marking..

Turning to group effects, although we did not observe a significant *group* by *antecedent* interaction as expected ($H_{5.4}$), we did observe a significant three-way interaction with *order*. This surfaced as a greater differential effect of *order* over antecedent in the experimental group (see Figure 5.10). In particular, it appears[68] that the experimental group exhibited a stronger effect of antecedent at the beginning of the experiment with both groups becoming increasingly insensitive to the type of antecedent by the end of the experiment. This was unexpected and is clearly incompatible with previous accounts for attrition couched in terms of features (Tsimpli et al. 2004) or overt pronouns as a processing default (Sorace 2011). Moreover, our results are inconsistent with the findings in Chamorro et al. (2015a). Recall they reported a weakening of the non-SpecIP bias in their experimental group, whereas we observed an apparent strengthening of the same bias. How to reconcile our findings with those in Chamorro et al. (2015a) is not immediately clear. However, the present online attrition results fit with the general pattern reported in chapter 3. In both the ambiguous pronominal items and the ambiguous PR/RC items, when attrition effects were observed, the experimental group appeared to rely more heavily on their interpretive biases than the control group did. As such, we take the increased non-SpecIP bias in the present experimental group to be non-spurious.

Under that assumption, one possibility to integrate our findings with those reported by Chamorro et al. (2015a) is that we might suggest that different languages studied played a role. To see how this might affect the data, recall that there are microparametric differences between Italian and Spanish. Although both are null-subject languages, some previous work directly comparing biases in native speakers of both languages has argued that the non-SpecIP bias for overt pronouns is weaker in Spanish than in Italian (e.g., Filiaci et al. 2014, Contemori & Di Domenico 2021). In fact, Filiaci et al. (2014) found that while their native Italian speakers exhibited a clear SpecIP bias in items like (79a), when they presented translational equivalents to native (Iberian) Spanish speakers as in (79b), no statistical differences were observed in their reading times of SpecIP and non-SpecIP disambiguated items with overt pronouns.

(79) a. *Dopo che Giovanni ha criticato Bruno così ingiustamente,* | *(ø / lui) si è scusato*
after that Giovanni has criticised Bruno so unjustly | (pro / he) REFL is excused
ripetutamente. [Italian]
repeatedly

'After Giovanni embarrassed Bruno so unjustly, he apologised repeatedly.'

b. *Después de que Bernardo criticó a Carlos tan injustamente,* | *(ø / el) le pidió*
after of that Bernardo criticised DOM Carlos so unjustly | (pro / he) CL asked.for
disculpas. [Spanish]
forgiveness

'After Bernardo criticised Carlos so unjustly, he apologised.'

Admittedly, they tested items with the order embedded-matrix, whereas both the present study and Chamorro et al. (2015a) used the order matrix-embedded. Nonetheless, if we suppose that native (Iberian) Spanish speakers are already relatively more flexible in their interpretation of overt pronouns with regard to the PAH, we might speculate that under attrition this is not a reliable enough bias, and as a result, speakers

[68] This interpretation is based on visual inspection. Given the continuous nature of *order*, *post hoc*s like those used in Aguilar et al. (2021) would require us to arbitrarily divide up the predictor resulting in a loss of information. Nonetheless, for completeness, Table D.3 presents simplified models over the first and second half of items separately, following Aguilar et al. (2021). Those models indicated a significant *group* by *antecedent* interaction only in the first half of the experiment.

become yet more flexible. Conversely, in Italian where the SpecIP bias has been argued to be stronger, it may be that speakers undergoing attrition are more able to rely on this structural bias when faced with local ambiguities.

However, such an interpretation runs into two related empirical problems and should be rejected. First, recall that in three distinct late-L2 samples of native (Iberian) Spanish speakers, Martín-Villena (2023) found an increased non-SpecIP bias relative to functional monolinguals in the interpretation of globally-ambiguous items like (80). Second, in the Italian version of the same items (82), Tsimpli et al. (2004) found that native Italian speakers undergoing attrition exhibited a weaker – not stronger – non-SpecIP bias for overt pronouns. That is to say, there is evidence that attrition can manifest as either a strengthening or a weakening of the non-SpecIP bias for overt pronouns in both Italian and Spanish. This is incompatible with any potential explanation for these inconsistencies in terms of microparametric differences between the two languages.

(80) *La anciana saludó a la mujer cuando ella cruzaba la calle.* [Spanish]
the elderly.woman greeted DOM the woman when she crossed the street

'The elderly woman greeted the woman when she crossed the street.'

It is also unclear how we could attribute the differences between the finding in the above-mentioned studies to the participant samples. Recall from section 5.5.1.1, that the present experimental group had been living in their L2 for a prolonged period (mean = 12.14 years), had not travelled back to their L1 community in the two months prior to testing, frequently used their L2 English (mean = 82.58% of a typical day), infrequently used their L1 (mean = 16.16%), and were upper-intermediate to advanced L2 speakers of English as measured by the Cambridge Assessment General English quick placement test. All of this not only suggests that this is an ideal group in which to explore attrition effects, but also that this group is highly similar to that tested in Chamorro et al. (2015a).

In a similar vein, it also seems implausible that we could explain the conflicting findings via the experimental manipulations or methodologies employed. To see why this is the case, remember that the present study used a gender manipulation whereas Chamorro et al. (2015a) used a number manipulation. This is potentially relevant as Carminati (2005) found that the penalty for violating the PAH (in Italian) was smaller for null pronouns forced toward a non-SpecIP antecedent when they were disambiguated via gender (81a) than when they were disambiguated by number (81b). If we assume that a similar asymmetry holds for overt pronouns which were not tested by Carminati (2005), then we might try to revive the idea that the baseline non-SpecIP bias was weaker in Chamorro et al. (2015a), resulting in the different attrition effects. However, this line of reasoning faces the same empirical difficulties as the 'microparametric' idea discussed above. Namely, Martín-Villena (2023) observed a strengthening of the non-SpecIP bias for overt pronouns whereas Tsimpli et al. (2004) found a weakening despite Martín-Villena (2023) explicitly designing their sentence interpretation task to be as similar as possible to the one used in Tsimpli et al. (2004) and using minimally adapted items. This would remain unaccounted for if we were to reduce the differences between our results and those in Chamorro et al. (2015a) to the differences in the experimental manipulation (i.e., gender vs number) or methodology (i.e., self-paced reading vs eye-tracking-while-reading). Thus based on the available data, it appears that we are forced to acknowledge that attrition may lead to opposing outcomes for the interpretation and processing of overt pronouns in Italian and Spanish with no clear indication of what drives this inconsistency.

(81) a. *Quando Maria lo cerca, | ø diventa ansios-o.* [Italian]
when Maria.F CL.M looks.for *pro* becomes anxious-M

'When Maria looks for him, he becomes anxious.'

b. *Quando i Rossi lo cercano,* | *ø diventa ansios-o.*
when the Rossi.PL CL.SG looks.for *pro* becomes anxious-SG

'When the Rossis looks for him, he becomes anxious.'

Turning finally to the results for the reading times of the disambiguating window for items containing null pronouns, participants exhibited no bias (i.e., SpecIP and non-SpecIP disambiguated items were read similarly) as predicted ($H_{5.3}$) based on the results of Chamorro et al. (2015a). This raises the question of how we should integrate this finding with our findings for the overt pronouns in the same experiment or the findings for null pronouns in chapter 3. Recall that there we looked at the interpretation of ambiguous items in the matrix-embedded order as in (82). In that experiment, results indicated that the experimental group (Italians living abroad) exhibited a stronger interpretive bias than the control group. In particular, while the control group interpreted the null subject as referring to the matrix subject around chance level (46.97%), the experimental group selected subject interpretations noticeably more frequently (62.22%) replicating the results reported in Tsimpli et al. (2004). Thus attrition can replicably result in an increased SpecIP bias in the interpretation of null pronouns.

(82) *L'anziana signora saluta la ragazza quando (ø / lei) attraversa la strada.* [Italian]
The.elderly woman greets the girl when (*pro* / she) crosses the road

'The elderly woman greets the girl when she crosses the road.'

Additionally, in the self-paced reading experiment presently under discussion, the experimental participants exhibited an increased non-SpecIP bias in their online processing of overt pronouns (at least toward the beginning of the experiment). If that were driven by a general tendency to rely more heavily on one's L1 biases under attrition as we suggested, then the absence of a similar effect for the same participants' processing of null pronouns would be unexpected. As such, based on the results for the null pronouns, we might argue that the attrition effect within the overt pronouns should not be accounted for in terms of a general increased reliance on one's L1 biases.

As an initial response to this, one might be to point out the differences between the gender manipulations within the null and over items. In the null items, the disambiguating gender marking was only present on the embedded secondary predicate (e.g., affamat-o_M/-a_F – 'hungry') whereas gender marking was also manifest on the pronouns in overt items. Thus, we might wonder if the gender marking on the embedded predicates was 'insufficiently' salient to disambiguate the null items, resulting in the observed asymmetry. This can be straightforwardly dismissed in two ways. First, participants' responses to comprehension questions were at ceiling for the null items, indicating that they were attending to the gender marking on the embedded predicates at least in their final interpretations. Second, preempting our discussion of the PR/RC items in chapter 6, participants reliably exhibited differences in their reading times when we used gender-marking on the embedded secondary predicate to force readings with or against their preferred parses. We take this to indicate that although the gender marking on the secondary predicates may potentially be less salient, participants were still sensitive to it in online processing.

Nonetheless, there are two further reasons that we should still be cautious of taking the asymmetry between our null and overt pronouns as evidence against our 'increased biases' interpretation. First, it rests on a null result. Second, we have independent reasons to believe that our items may not have been suitable to detect the SpecIP bias for null pronouns given the results from the two experiments with native speakers still living in Italy. Hence, in line with what we argued for in the discussion of the comprehension responses, we refrain from interpreting the lack of a group effect. Nonetheless, it may be useful for future work to investigate attrition to further explore this potential asymmetry using items in which native Italian speakers have been reliably found to exhibit a SpecIP bias (e.g., the items from Carminati 2002 borrowed by Filiaci et al. 2014).

5.7 Chapter summary

In this chapter, we presented three experiments with locally ambiguous pronominal sentences. These experiments were conducted to investigate online PAH biases both in native Italian speakers living in Italy and in potential attriters living abroad. Overall, results from the two experiments with native speakers still living in Italy were very similar, indicating no bias for null pronouns and a non-SpecIP bias for overt pronouns in our items. Comparison of the experimental group (Experiment 3) and the control group (Experiment 1) for our self-paced reading task indicated that attrition affected overt pronouns but not null pronouns as expected. However, the effect appeared in the unpredicted direction. *Contra* what was reported by Chamorro et al. (2015a), attrition appeared as a strengthening of the non-SpecIP bias for overt pronouns and was subject to adaptation effects.

6 Attrition and the PR-First Hypothesis in locally ambiguous items

6.1 Introduction

In this chapter, we present three experiments to investigate the online resolution of PR/RC parsing ambiguities in Italian, looking at native speakers living both in Italy and abroad. In section 6.3, we first present the results from a self-paced reading experiment with native Italian speakers residing in Italy. This experiment used the locally-ambiguous PR/RC items reported in chapter 4. Given the relatively limited previous work on the online parsing of PRs, this experiment was conducted to investigate the online predictions drawn from the PR-First Hypothesis. As predicted, results indicated a LA bias when PRs were locally blocked as well as a HA bias when PRs were locally available. Results also indicated a persistence effect of PR parses in items that were locally compatible with PR parses but then later disambiguated to force a RC reading. To follow up on that effect, section 6.4 presents a sister experiment which used an eye-tracking-while-reading task. For that second experiment, the items were identical to those in the self-paced reading experiment and participants again consisted of native Italian speakers still residing in Italy. Overall, results from the second experiment converge on those from the first, in particular indicating that the persistence effect with PRs is replicable. As such, we take this to be a new type of evidence in support of the PR-First Hypothesis. To explore how attrition affects PR/RC parsing biases, section 6.5 presents a third experiment in which we presented native Italian speakers living in a majority English-speaking country with the same self-paced reading task as in section 6.3. Results indicated that the Italians living abroad resolved these ambiguities in a similar way to Italians living in their home community. As such, our results did not replicate the change in parsing biases reported in Dussias (2003, 2004) and Dussias & Sagarra (2007).

6.2 Research questions

In this chapter, we focus on the following research questions:

(RQ_6) In native Italian speakers still living in Italy, can we observe a locality effect in the processing of items that unambiguously contain RCs?

(RQ_7) In native Italian speakers still living in Italy, can we observe a PR-firstness effect in the processing of items that are PR-compatible?

(RQ_8) If so, how are these biases affected by attrition?

6.3 Experiment 1: Self-paced reading

6.3.1 Method

6.3.1.1 Participants

For this self-paced reading experiment, participants consisted of 66 native Italian speakers (Female = 34; Male = 32)[69] who (i) had lived in Italy from birth until at least the age of 16, (ii) were living in Italy at the time of testing, (iii) had no diagnosed language related disorder, and (iv) reported growing up monolingually.[70] At the time of testing, the average age of participants was 43.47 years (SD = 7.39 years).[71] As in our earlier experiments, participants were asked to report how many languages they could speak at or above an 'intermediate' level. They were also asked to indicate how many hours they used each of these languages in a typical day. Of the 66 participants, all but 2 indicated that they could speak another language (and all but 4 could speak English). To quantify participants' language use, we calculated their percentages for each language (i.e., hours of language x / total hours reported for any language). Table 6.1 reports the means and standard deviations for their use of Italian, English, and 'Other.'

	Mean	SD
Italian	82.85%	14.28%
English	13.07%	11.84%
Other	4.08%	6.88%

Table 6.1: Mean language use by participant as percentages of a typical day with standard deviations.

6.3.1.2 Stimuli

The final 32 PR/RC items from the norming study (chapter 4) were used as the critical items. Recall that each contained a complex NP (NP_1 of NP_2) followed by a finite embedded CP. Within these items, we manipulated PR availability and attachment. To manipulate PR availability, we exploited the selectional properties of the matrix predicate (counterbalanced *within* items); half of the time items appeared with a matrix predicate of perception and half the time they appeared with a non-perceptual predicate. To force attachment, we used gender marking on a secondary predicate within the embedded CP (-o.M, -a.F). For that reason, the two NPs in the complex NP always differed in gender (with the gender of NP_1 counterbalanced *across* items). An example item in all 4 conditions is provided in (83) with vertical bars to indicate segmentation.

(83) a. Perceptual HA [Italian]

Gianni | ha visto | il collega | della biologa | che correva | sporco | di fango.

[69] These are slightly different participants from those reported in section 5.3. Due to high rates of data loss for some participants in the relevant sets of items (> 50%), we could not consider all 68 of the original participants for both the pronominal and PR/RC items. As such, we do not consider 2 of the original participants in the present analyses as well as a different 2 participants in the pronominal analysis.

[70] Despite all participants reporting growing up monolingually at the pre-screening, when specifically asked about potential exposure to 'dialect' as a child (for example with one's grandparents), 14 participants reported some exposure.

[71] As mentioned in footnote 42, although our participants are somewhat older than those in Chamorro et al. (2015a) or Martín-Villena (2023), they are more comparable to those in Tsimpli et al. (2004: Ianthi Tsimpli, p.c.). Moreover, in each experiment, our experimental and control groups are comparable in age. As such, age is highly unlikely to drive any potential group-level effects unlike what was reported in Kaltsa et al. (2015) who reported a 30 year age difference between groups.

b. Perceptual LA

Gianni | *ha visto* | *il collega* | *della biologa* | *che correva* | *sporca* | *di fango.*
Gianni has seen the colleague.M of.the biologist.F that ran dirty.F of mud

'Gianni saw the colleague of the biologist that was running covered in mud.'

c. Non-perceptual HA

Gianni | *vive con* | *il collega* | *della biologa* | *che correva* | *sporco* | *di fango.*

d. Non-perceptual LA

Gianni | *vive con* | *il collega* | *della biologa* | *che correva* | *sporca* | *di fango.*
Gianni lives with the colleague.M of.the biologist.F that ran dirty.F of mud

'Gianni lives with the colleague of the biologist that was running covered in mud.'

Sentences for the critical PR/RC items were distributed across 4 lists such that participants saw each item only once, with 8 sentences per condition. The experiment also included unambiguous biclausal fillers (N = 40) and the 32 pronominal items reported in chapter 5 for a total of 104 items.

6.3.1.3 Procedure

The procedure for this experiment was the same as for the pronominal items (section 5.3.1.3). To re-iterate, we first obtained ethical approval from Ethics Committee of the Faculty of Modern and Medieval Languages and Linguistics at the University of Cambridge. We then recruited participants through Prolific Academic and ran the experiment via PCIbex (Zehr & Schwarz 2018). After providing informed consent, participants filled out a short language history questionnaire. This was followed by the self-paced reading task with a moving window (window boundaries indicated by '|'s in 83). After reading each item, participants were presented with a polar comprehension question. For the critical items, this always asked about the subject of the non-matrix CP (with the questioned NP counterbalanced *within* items) and the possible answers were labelled 'F' (yes) and 'J' (no). To respond, participants pressed the relevant key causing the experiment to pass to the next item.

Recall that we included these comprehension questions as a way to tap persistence of an initial/preferred parse. This was because previous studies on different types of garden-path sentences have found that readers' initial/preferred parse may persist even after reanalysis/re-ranking (e.g., Christianson et al. 2001, Sturt 2007, Slattery et al. 2013). This can be exploited to derive a secondary and indirect way of tapping parser biases, namely by comparing response accuracy after items in which we expect reanalysis and after items in which we do not. If accuracy is lower for the items that are expected to require reanalysis, this would support the idea that reanalysis has indeed taken place.

6.3.1.4 Data cleaning and analysis

In the following, we present two planned analyses as well as one unplanned analysis. For the planned analyses, we consider the reading times of the disambiguating window (which consisted of the secondary depictive predicate) as well as the accuracy to the comprehension questions. For the unplanned analysis, we also consider the reading times of the post-critical window (which consisted of the sentence-final PPs). We opted to present this additional analysis given De Santo & Lee's (2022) additional analysis of the data reported in Lee & De Santo (2022). Despite learning of their reanalysis after beginning data collection, all of our items included sentence final PPs to push the critical region off the sentence edge allowing us to also perform the additional analysis.

Due to a coding error, we lost 3 items for 12 participants. To clean the remaining reading time data, we coded as missing any trial for which the analysed window, or any window prior to it, had an implausibly fast reading time ($<$ 200 ms). This affected 1.83% of trials in both the critical and post-critical windows. For both windows, we then calculated the IQR for each condition and coded as missing any value that lay 1.5 IQR above the third quartile for that window. This affected a further 6.79% of trials in the critical window (for a total data loss of 8.62%) and 7.27% of trials in the post-critical window (total data loss: 9.10%). The remaining data was equally distributed across the 4 conditions (critical: $\chi^2(1,1897) = 0.23$; $p = 0.93$; post-critical: $\chi^2(1,1887) = 0.98$; $p = 0.81$). For the analysis of the comprehension questions, we only consider those items included in the planned analysis of the critical window.

To analyse the data, we used the *lme4* package (Bates et al. 2015) in R (R Core Team 2022). We started by identifying the best random-effects structure. To that end, we conducted a family of intercept-only models in which we considered the various random-effects structures of our theoretically relevant fixed predictors (i.e., random intercepts by *item* and *participant* as well as random slopes by *attachment* and *predicate*). From these, we selected the best fitting random effects structure (Matuschek et al. 2017) using the AIC. We then fit a base model with only the fixed effects that we were theoretically interested in using sum coding (i.e., -0.5, 0.5). For *attachment*, we coded LA as the negative level. For *predicate*, we coded non-perceptual as the negative level. As previous work has suggested that trial order may obfuscate one of the effects we are interested in (i.e., PR-firstness) due to adaptation to the experimental context, we then conducted a further two models which included *order* (centred over the experiment) as either a simple effect or potential interaction term. We compared these additional two models against the base model using the AIC to assess whether the added terms improved the fit. If complicating the model reduced the AIC by 2 or more, the more complicated fixed effect structure was maintained. Otherwise, we selected the simpler model. For analysis of the accuracy to comprehension questions, once we had identified the best fitting model with regard to *order* we conducted two further models in the same way with *response target* (negative level: yes). This was done to account for any potential effects of the acquiescence bias (Holbrook 2008).

6.3.1.5 Hypotheses and predictions

The following hypotheses were drawn:

($H_{6.1}$) In the non-perceptual items, participants will exhibit a LA bias.

($H_{6.2}$) In the perceptual items, participants will exhibit a HA bias.

For reading times, we predict that LA disambiguated items should be read faster than HA disambiguated items when the matrix predicate is non-perceptual due to locality (e.g. Late Closure). When the matrix predicate is perceptual, however, we expect the opposite. HA items should be read faster than LA items due to the PR-First Hypothesis. This should result in a two-way interaction of *attachment* and *predicate*. If the best fitting model includes also a three-way interaction with *order*, this would be expected to surface as a PR-firstness effect in the first half of the experiment that disappears in the second half due to adaptation (cf., Pozniak et al. 2019, Aguilar et al. 2021).

As we take response accuracy in this context to be an indirect measure of processing biases, we expect the same pattern of results. When the matrix predicate was non-perceptual we expected greater accuracy for LA items than HA ones. Conversely, when the matrix predicate is perceptual, there should be greater accuracy with HA items than LA ones resulting in an interaction of *attachment* and *predicate*. Should there be a three-way interaction this is expected to be due to a reduced PR-firstness effect as the experiment progresses. No specific predictions were made regarding *response target*.

6.3.2 Results

6.3.2.1 Reading times in the critical window

Figure 6.1 presents the mean reading times for the four conditions in the critical window. For the non-perceptual items, participants read LA-disambiguated items (mean = 700.47 ms) faster than HA-disambiguated ones (mean = 811.79 ms). For perceptual items, this pattern was reversed and participants read LA disambiguated items (mean = 763.24 ms) slower than HA disambiguated items (mean = 718.97 ms).

Reading times were logged and subjected to a family of mixed-effect regressions as described above. Model comparison indicated that the best fitting model included *order* as a simple effect (improvement to AIC $>$ 2) but not as a potential interaction term. The output for that model is reported in Table 6.2. That model indicated a significant effect of *order* ($\hat{\beta}$ = -0.001; t = -4.60; $p < 0.001$) indicating that participants read more quickly as the experiment progressed. The model did not indicate significant effects of *predicate* or *attachment*, but there was a significant interaction between the two ($\hat{\beta}$ = -0.15; t = -4.56; $p < 0.001$). To follow up on this interaction, we ran pairwise comparisons of *attachment* within *predicate* using the *emmeans* package Lenth (2022). For these comparisons, we present Holm-Bonferroni corrected p-values. Results indicated a significant effect of *attachment* within the non-perceptual items ($\hat{\beta}$ = 0.10; t = 4.42; $p < 0.001$) with LA disambiguated items being read more quickly. Results also indicated a significant effect within the perceptual items ($\hat{\beta}$ = -0.05; t = -2.02; p = 0.04). In those items, however, this surfaced as faster reading times for HA disambiguations.

	Estimate	Std. Error	t-value	p-value	
Intercept	6.52	0.04	165.19	$<$ **0.001**	***
Predicate	-0.01	0.02	-0.86	0.39	
Attachment	0.03	0.02	1.71	0.09	·
Order	-0.001	0.0003	-4.60	$<$ **0.001**	***
Predicate:Attachment	-0.15	0.03	-4.56	$<$ **0.001**	***

Table 6.2: Model output for the reading times of the critical window.

6.3.2.2 Reading times in the post-critical window

Figure 6.2 presents the mean reading times for the four conditions in the critical window. For the non-perceptual items, participants read LA-disambiguated items (mean = 777.49 ms) faster than HA-disambiguated ones (mean = 971.94 ms). For perceptual items, participants read LA disambiguated items (mean = 963.56 ms) slower than HA disambiguated items (mean = 854.33 ms).

Reading times were logged and subjected to a family of mixed-effect regressions. Model comparison indicated that the best fitting model included *order* as a simple effect, but not as a potential interaction term. The output for that model is presented in Table 6.3. The model indicated a significant effect of *attachment* ($\hat{\beta}$ = 0.05; t = 2.66; $p < 0.01$) indicating that overall participants read HA disambiguated items more slowly. There was also a significant effect of *order* ($\hat{\beta}$ = -0.003; t = -9.72; $p < 0.001$) with participants reading more quickly as the experiment progressed. The effect of *predicate* was not significant but its interaction with *attachment* was ($\hat{\beta}$ = -0.30; t = -8.13; $p < 0.001$). To follow up on that interaction we again conducted pairwise comparisons of *attachment* within *predicate* using the *emmeans* package (Lenth 2022). For these, we report Holm-Bonferroni corrected p-values. The pairwise comparisons indicated a significant effect of *attachment* within the non-perceptual items ($\hat{\beta}$ = 0.20; t = 7.55; $p < 0.001$) with LA disambiguated items

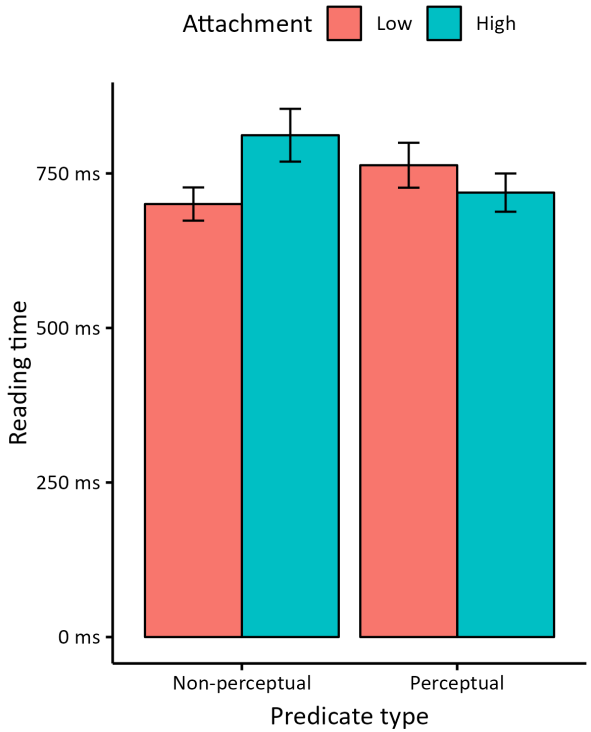

Figure 6.1: Global average reading times for the critical window with 95% confidence intervals.

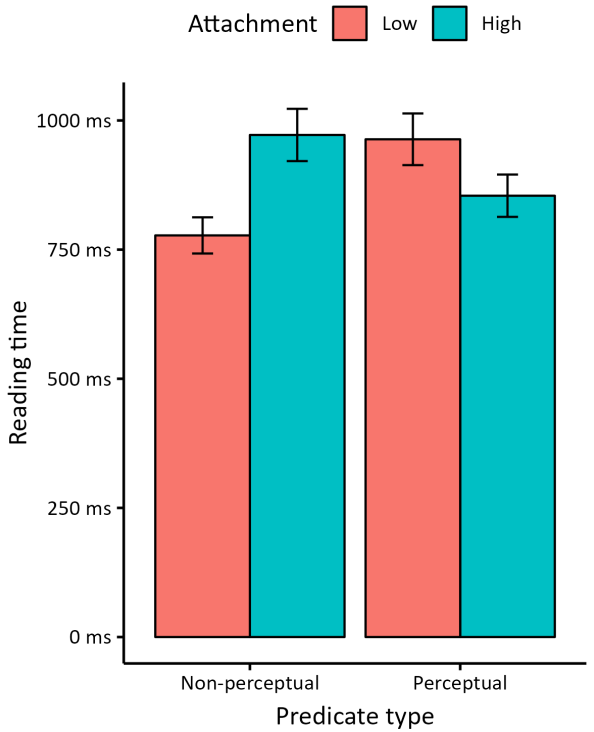

Figure 6.2: Global average reading times for the post-critical window with 95% confidence intervals.

being read more quickly. There was also a significant effect of *attachment* within the perceptual items ($\hat{\beta}$ = -0.10 t = -3.90; $p < 0.001$) with LA disambiguated items being read more slowly.

	Estimate	Std. Error	t-value	p-value	
Intercept	6.67	0.04	167.84	$<$ **0.001**	***
Predicate	0.02	0.02	1.17	0.24	
Attachment	0.05	0.02	2.66	$<$**0.01**	**
Order	-0.003	0.0003	-9.72	$<$ **0.001**	***
Predicate:Attachment	-0.30	0.04	-8.13	$<$ **0.001**	***

Table 6.3: Model output for the reading times of the post-critical window.

6.3.2.3 Response accuracy for comprehension questions

Figure 6.3 presents the mean accuracy (by participant) for the four conditions. In the non-perceptual items, participants gave the correct answer at similar rates for both LA (mean = 91.98%) and HA disambiguated items (mean = 89.16%). In the perceptual items however, participants' responses were noticeably more accurate after HA disambiguated items (mean = 95.45%) than LA disambiguated ones (mean = 83.76%).

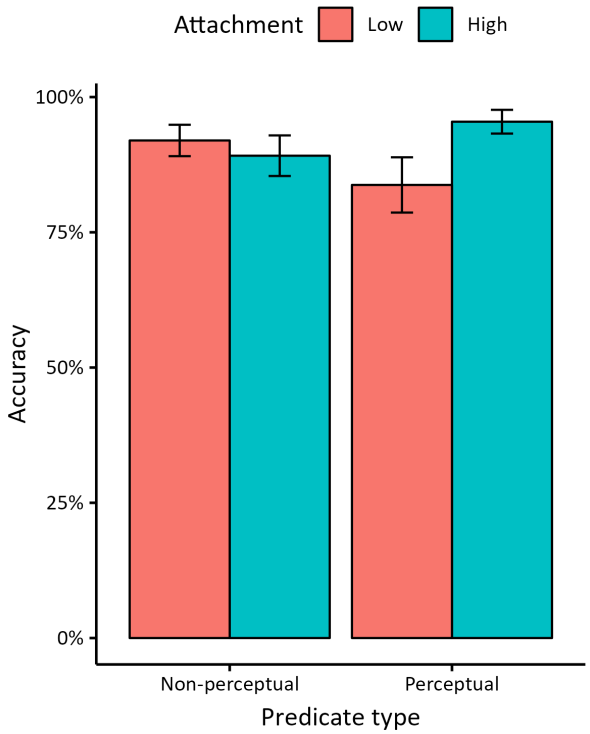

Figure 6.3: Average response accuracy (by participant) with 95% confidence intervals.

Responses were coded as ± correct and entered into a family of mixed-effect logistic regressions as described above. Model comparison indicated that the best-fitting model did not include *order*, but did include *response target* as a simple effect. The output for that model is presented in Table 6.4. The model indicated a significant effect of *response target* ($\hat{\beta}$ = -0.46; z = -2.67; $p < 0.01$) with participants giving more accurate responses when the target response was *sì* ('yes'). The effects of *predicate* and *antecedent* were not significant but their interaction was ($\hat{\beta}$ = 2.04; z = 5.61; $p < 0.001$). To follow up on that interaction

	Estimate	Std. Error	*z*-value	*p*-value	
Intercept	2.94	0.21	13.89	$< \mathbf{0.001}$	***
Predicate	0.06	0.18	0.35	0.73	
Attachment	0.61	0.37	0.167	0.10	·
Response target	-0.46	0.17	-2.67	$< \mathbf{0.01}$	**
Predicate:Attachment	2.04	0.36	5.61	$< \mathbf{0.001}$	***

Table 6.4: Model output for response accuracy.

we conducted pairwise comparisons of *attachment* within *predicate* and report Holm-Bonferroni corrected *p*-values as above. That indicated that the effect of *attachment* was non-significant within the non-perceptual items ($\hat{\beta}$ = -0.41; z = -1.02; p = 0.31). Within the perceptual items however, the effect of *attachment* was significant ($\hat{\beta}$ = 1.63; z = 3.88; $p < 0.001$) with participants giving more accurate responses after HA items.

Given the relatively small number of incorrect responses (N = 194, 10.23% of responses) as well as relatively small differences in error rates for the comprehension questions (just under 12 percentage points in the perceptual items), we additionally inspected which items participants got wrong. This was done to see if the trend reported above could have been driven by a small set of potentially bad items. Inspection of the items however indicated that errors were fairly well distributed with all items receiving some wrong responses (mean accuracy by item = 89.77%; SD = 4.06%). To look for potential outliers, we calculated the IQR for accuracy by item. For no items, did accuracy fall 1.5 IQR below/above the 1st/3rd quartile.

6.3.3 Interim discussion

In the proceeding, we have presented a self-paced reading task to investigate the online parsing predictions of the PR-First Hypothesis. Specifically, we were interested in whether Italian speakers exhibit a LA bias with true RCs as well as a HA bias when PRs are locally available.

For items with a non-perceptual matrix predicate, reading times for both the critical and post-critical windows were significantly longer when the RC was forced to attach high than when it was forced to attach low. Thus despite Italian being a so-called 'HA language,' a LA bias emerges in the online processing of true RCs once PRs are locally blocked. This was as predicted ($H_{6.1}$) and is convergent with Lee & De Santo's (2022) and Aguilar et al.'s (2021) results from Italian and Spanish respectively.

When the matrix predicate was perceptual, however, reading times for both the critical and post-critical windows were significantly longer when the embedded CP was forced to attach low, confirming ($H_{6.2}$). As LA is incompatible with PRs in our items, we interpret this as an online bias for PRs where locally available. Again this is convergent with the previously reported results, although this effect was not restricted to the post-critical window (cf. De Santo & Lee 2022) nor we did find evidence for a loss of PR-firstness due to adaptation to the experimental context (cf. Aguilar et al. 2021).

Interestingly, the online bias for PRs also appears to have affected offline comprehension responses. That is to say, we observed greater accuracy after HA perceptual items than LA ones. This is what we would expect if (i) participants first projected a PR parse in those items upon encountering the embedded CP and (ii) that intermediate parse (which must be discarded upon encountering a LA disambiguation) may persist in short term memory and interfere when responding to the comprehension questions. However, we may want to be sceptical of this finding and interpretation given that we did not find a persistence effect in the non-perceptual items despite a clear LA bias in the reading times of both windows. As such, before trying to interpret this asymmetry, it is worth checking if the result is replicable. To that end, we now turn to a sister experiment

using eye-tracking-while-reading.

6.4 Experiment 2: Eye-tracking-while-reading

6.4.1 Method

6.4.1.1 Participants

Participants for this experiment were the same 27 native speakers from section 5.4. To re-iterate, all participants (i) had lived in Italy from birth until at least the age of 16, (ii) were living in Italy at the time of testing, (iii) had no diagnosed language-related disorder, and (iv) did not report speaking any other languages at home as a child apart from some *dialetto* ('dialect').[72] At the time of testing, participants had an average age of 41.37 years (SD = 13.34). Table 6.5 re-reports the group's use of Italian, English and 'Other' as percentages of a typical day.

	Mean	SD
Italian	79.35%	16.52%
English	12.67%	13.38%
Other	7.98%	11.30%

Table 6.5: Mean language use by participant as percentages of a typical day with standard deviations, repeated from Table 5.5.

6.4.1.2 Stimuli

The stimuli were identical to those in the self-paced reading experiment and interest areas were defined based on the windows of the same experiment. Examples are provided in (84) with vertical bars to indicate the critical interest area.

(84) a. Perceptual HA [Italian]

Gianni ha visto il collega della biologa che correva | sporco | di fango.

b. Perceptual LA

Gianni ha visto il collega della biologa che correva | sporca | di fango.
Gianni has seen the colleague.M of.the biologist.F that ran dirty.F of mud

'Gianni saw the colleague of the biologist that was running covered in mud.'

c. Non-perceptual HA

Gianni vive con il collega della biologa che correva | sporco | di fango.

d. Non-perceptual LA

Gianni vive con il collega della biologa che correva | sporca | di fango.
Gianni lives with the colleague.M of.the biologist.F that ran dirty.F of mud

'Gianni lives with the colleague of the biologist that was running covered in mud.'

[72] When specifically asked about potential exposure to 'dialect' as a child (for example with one's grandparents) seven reported some exposure.

6.4.1.3 Procedure

The procedure was the same as that reported in section 5.4.1.3. To summarise, we used an EyeLink 1000 Plus (SR Research) and stimuli were presented at a distance of 87 cm on a 48 by 27 cm screen. We recorded the movements of participants' right eyes with a 35mm lens at a sampling rate of 1000 Hz. Items always appeared as a single (left-aligned) line of text in 32-point font (Calibri). Participants were tested individually in a sound-attenuated room in Italy.

After reading the instructions, participants performed a 9-point calibration, followed by a practice phase. Before each item, we used a cross on the left-hand side of the screen to perform a drift check. If this were successful, the item appeared immediately to the right of the cross. If this the drift check were unsuccessful, it would initiate a re-calibration sequence. After reading the item, participants were presented a comprehension question. The comprehension question and possible answers (labelled 'F' and 'J') were presented alone on the screen. To respond, participants pressed the corresponding key.

6.4.1.4 Data cleaning and analysis

In the following, we conduct two planned analyses, one over the eye-tracking measures for the critical interest area and one over the accuracy to the comprehension questions. For consistency with the analysis of the self-paced reading task reported in section 6.3 we also present an unplanned analysis over the eye-tracking measures for the post-critical interest area.

To filter our eye-movement data, we first tried to merge any fixations under 80 ms with another fixation within 0.5 degree of visual angle. Where this was not possible, the fixations were removed. We also filtered out any fixation greater than 1000 ms. From the fixation data, we then extracted First Pass, Go Past, and Total Time measures. The First Pass (FP) is the sum of any fixations from the first time the region is entered from the left until the region is first exited (either to the left or the right). The Go Past (GP) is the sum of any fixations from the first time the region is entered from the left until it is exited to the right. This can include fixations to the left of the region. Finally, the Total Time (TT) is the sum of all fixations within a region regardless of whether it's the initial reading or a re-reading.

For the analysis of the critical region, we first coded as missing any trials in which the critical region itself was not fixated. This affected 1.16% of the data. For the comprehension questions, we additionally excluded missing responses. This affected a further 0.93% of the PR/RC trials (total exclusion = 2.08%). For the FP and GP measures, we further coded as missing any trials for which the critical region was skipped in the initial reading (i.e., material to the right was read first). This affected a further 3.82% of the PR/RC data (total loss = 4.98%). Within each data set, the remaining data was equally distributed across the conditions (PR/RC: FP/GP: $\chi^2(1, 821) = 0.17$; $p = 0.98$; TT: $\chi^2(1, 854) = 0.005$; $p = 1.00$; Accuracy: $\chi^2(1, 846) = 0.06$; $p = 1.00$).

For the unplanned analysis of the post-critical region for the PR/RC items we first coded as missing any trials in which the critical region was not fixated before the post-critical region (4.98% data loss). We then coded as missing any trials in which the post-critical region itself was not fixated affecting a further 6.48% of the data for a total loss of 11.46%. The remaining observations were equally distributed across conditions ($\chi^2(1, 765) = 0.66$; $p = 0.88$) As the post-critical region was sentence-final, we did not consider GP times.

To analyse the data, we again conducted families of models to identify the best-fitting random-effect structure. We then fit a base model with only the fixed effects that we were theoretically interested in using sum coding (i.e., -0.5, 0.5). For *attachment*, we coded LA as the negative level. For *predicate*, we coded non-perceptual as the negative level. To account for potential adaptation effects, we then conducted a further two models which included *order* (centred over the experiment) as either a simple effect or potential interaction term. We then selected the best fitting model using the AIC. For analysis of the accuracy to comprehension questions, once we had identified the best fitting model with regard to *order* we repeated this process with

response target (negative level: yes).

6.4.1.5 Hypotheses and predictions

Our hypotheses were the same as in the self-paced reading experiment.

(H_1) In the non-perceptual items, participants will exhibit a LA bias.

(H_2) In the perceptual items, participants will exhibit a HA bias.

For eye-movement measures in the critical/post-critical region of the PR/RC items, we predict shorter durations for LA items than HA items when the matrix predicate is non-perceptual due to locality. When the matrix predicate is perceptual, however, we expect shorter durations for HA items than LA items due to the PR-First Hypothesis. This should lead to a two-way interaction of *attachment* and *predicate*. Should any of the best fitting models additionally include a three-way interaction with *order*, we expect this to surface as a weakening/loss of the PR-firstness effect as the experiment progresses due to adaptation to the experimental context.

For response accuracy, we expect greater accuracy for LA items than HA items when the matrix predicate is non-perceptual. When it is perceptual, however, we expect the opposite pattern. This should result in an interaction of *attachment* and *predicate*. If the best-fitting models include a three-way interaction with *order*, we would expect this to surface as a reduction of the PR-firstness effect over time.

6.4.2 Results

6.4.2.1 Eye-tracking measures in critical interest area

Table 6.6 presents the average FP, GP, and TT for the critical window. Figure 6.4 presents this same information graphically. We separately logged each measure and then subjected them to a family of 3 mixed-effect regressions. Model comparison indicated that for the FP and GP measures the inclusion of *order* did not improve the model fits (improvements to AIC < 2). For the TT measures, however, including *order* as a simple effect improved the fit (improvement to AIC > 2). The outputs for the best fitting models are presented in Table 6.7.

		FP		GP		TT	
Predicate Type	Attachment	Mean	CI	Mean	CI	Mean	CI
Non-perceptual	Low	337.28	23.17	631.40	86.31	804.03	72.90
	High	375.90	22.08	776.11	116.87	950.85	77.37
Perceptual	Low	377.29	26.16	771.19	129.71	918.48	79.52
	High	366.39	28.84	702.66	114.00	843.11	74.96

Table 6.6: Average FP, GP, and TT for the critical region with 95% confidence intervals.

In the TT model, there was a significant effect of *order* ($\hat{\beta}$ = -0.004; t = -6.72; $p < 0.001$). This surfaced as shorter TTs as the experiment progressed. None of the models indicated any effect of *predicate* or *attachment*. However, the FP and TT models both indicated a significant interaction of the two (FP: $\hat{\beta}$ = -0.13; t = -2.28; p = 0.02; TT: $\hat{\beta}$ = -0.28; t = -4.10; $p < 0.001$). The interaction was only marginal in the GP model ($\hat{\beta}$ = -0.19; t = -1.89; p = 0.06). To follow up on the significant interactions, we ran pairwise comparisons as above. In both the FP and TT measures, this indicated a significant effect of *attachment* within the non-perceptual items (FP: $\hat{\beta}$ = 0.11; t = 2.75; $p < 0.01$; TT: $\hat{\beta}$ = 0.19; t = 3.87; $p < 0.001$)

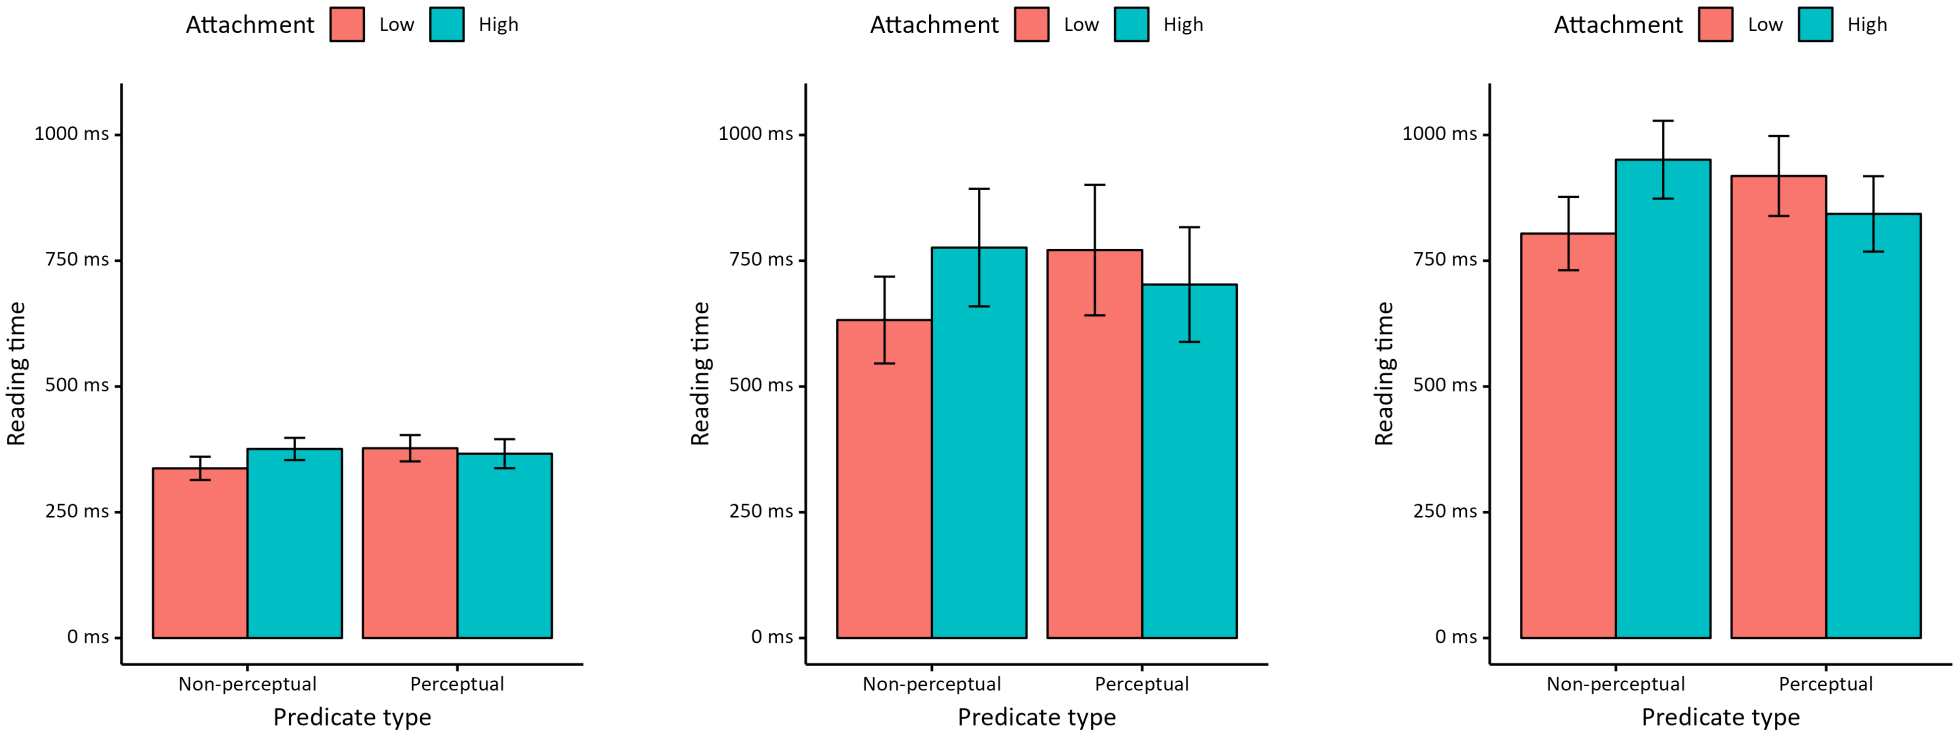

(a) FP (b) GP (c) TT

Figure 6.4: Eye-tracking measures for the critical region with 95% confidence intervals.

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	5.78	0.05	124.52	<**0.001**	***
Predicate	0.03	0.03	1.23	0.22	
Attachment	0.04	0.03	1.61	0.11	
Pred*Att	-0.13	0.06	-2.28	**0.02**	*

(a) FP

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.23	0.05	114.28	<**0.001**	***
Predicate	0.02	0.05	0.39	0.70	
Attachment	0.07	0.05	1.33	0.18	
Pred*Att	-0.19	0.10	-1.89	0.06	.

(b) GP

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.80	0.09	76.41	<**0.001**	***
Order	-0.004	0.0006	-6.72	<**0.001**	***
Predicate	-0.002	0.03	-0.07	0.94	
Attachment	0.05	0.03	1.39	0.16	
Pred*Att	-0.28	0.07	-4.10	<**0.001**	***

(c) TT

Table 6.7: Model outputs for the critical region.

which surfaced as shorter times for LA items. The effect of *attachment* within the perceptual items was not significant for either measure, but was marginal in the TT measure (FP: $\hat{\beta}$ = -0.02; t = -0.47; p = 0.64; TT: $\hat{\beta}$ = -0.09; t = -1.91; p = 0.06).

6.4.2.2 Eye-tracking measures in the post-critical interest area

Table 6.8 presents the average FP and TT reading times for the post-critical window. Figure 6.5 presents the same information graphically. We logged both measures and entered them into a family of regression as above. Model comparison indicated that the inclusion of *order* did not improve the fit for the FP measures (improvements to AIC $<$ 2), but the inclusion of *order* as a simple effect did improve the TT model fit (improvement to AIC $>$ 2). The outputs from the best fitting models are reported in Table 6.9.

		FP		TT	
Predicate Type	Attachment	Mean	CI	Mean	CI
Non-perceptual	Low	418.50	44.04	747.10	81.52
	High	472.26	54.87	838.35	103.01
Perceptual	Low	410.92	48.34	826.11	92.81
	High	452.35	50.88	755.17	92.39

Table 6.8: Average FP and TT for the postcritical region with 95% confidence intervals.

In the FP model, the model indicated a significant effect of *attachment* ($\hat{\beta}$ = 0.09; t = 2.37; p = 0.02) which surfaced as longer reading times for HA items. The effect of *predicate* and the interaction term were non-significant. In the TT model, we observed a significant effect of *order* ($\hat{\beta}$ = -0.003; t = -4.36; $p < 0.001$) with participants reading faster as the experiment progressed. That model did not indicate a significant effect of *attachment* or *predicate* and their interaction was only marginal.

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	5.78	0.09	62.57	<**0.001**	***
Predicate	-0.02	0.04	-0.44	0.66	
Attachment	0.09	0.04	2.37	**0.02**	*
Pred*Att	0.05	0.08	0.64	0.52	

(a) FP

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.47	0.13	51.33	<**0.001**	***
Order	-0.003	0.0007	-4.36	<**0.001**	***
Predicate	0.0006	0.04	0.02	0.99	
Attachment	0.005	0.04	0.13	0.89	
Pred*Att	-0.15	0.08	-1.92	0.06	.

(b) TT

Table 6.9: Model outputs for the post-critical region.

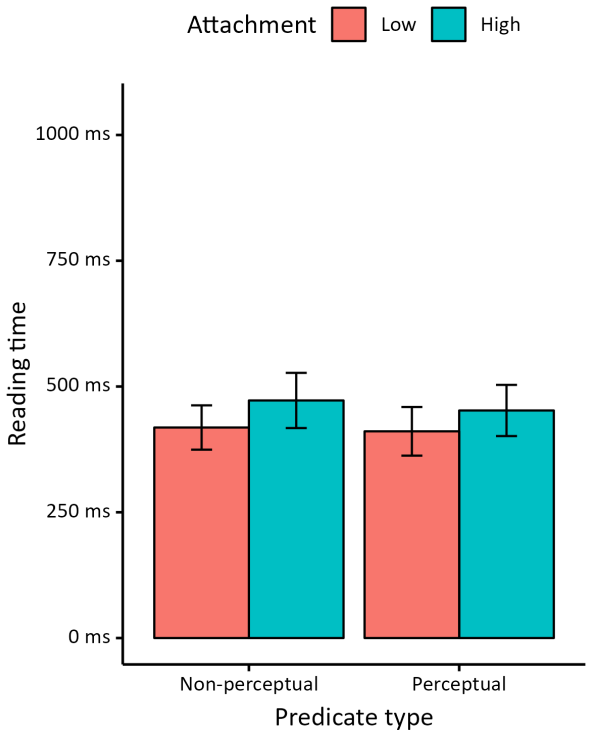

(a) FP

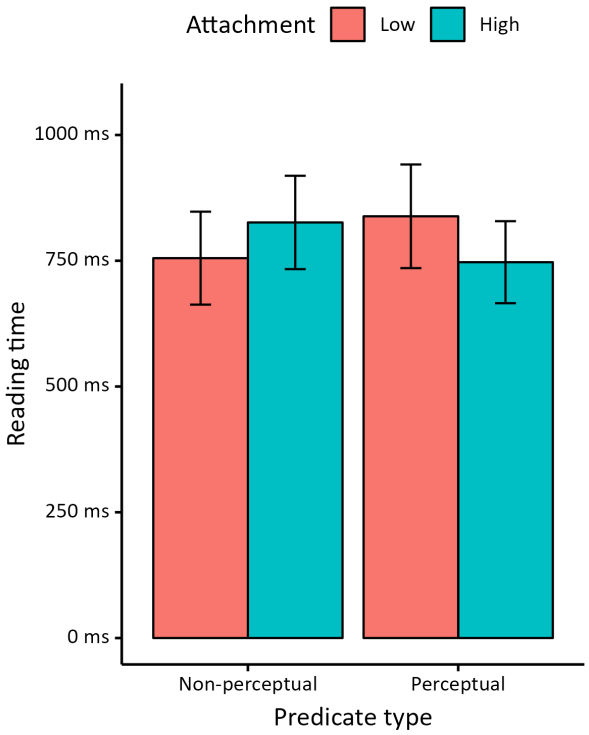

(b) TT

Figure 6.5: Eye-tracking measures for the post-critical region with 95% confidence intervals.

6.4.2.3 Response accuracy for comprehension questions

Figure 6.6 presents the average accuracy for the comprehension questions broken down by the matrix predicate type and attachment. In the non-perceptual condition, responses to LA and HA items had a similar accuracy (LA: 96.63%; HA: 96.76%). In the perceptual condition, however, the accuracy for LA items was noticeably lower than for HA items (LA: 90.74%; HA: 99.54%).

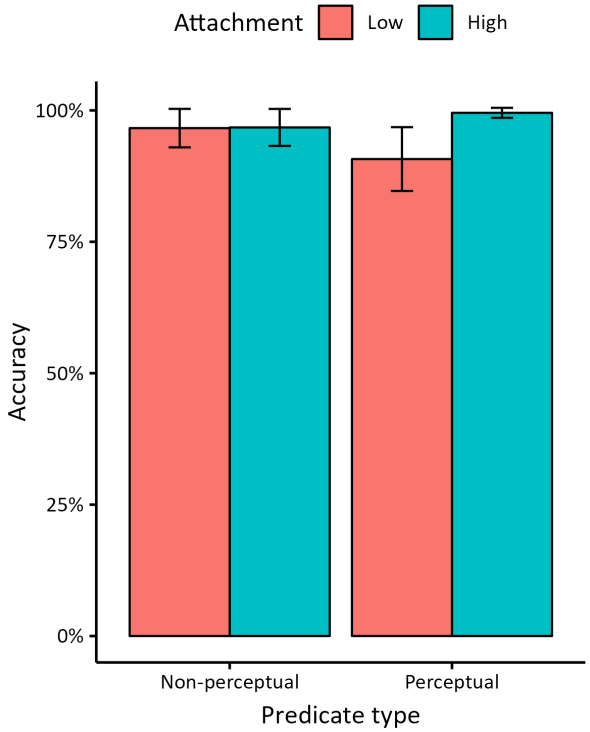

Figure 6.6: Average response accuracy (by participant) with 95% confidence intervals.

To analyse the accuracy for comprehension questions (coded as ± correct), we again conducted a family of logistic models. Model comparison indicated that the inclusion of *order* as a simple effect or potential interaction term did not improve the model fit (improvements to AIC < 2). However, the inclusion of *response target* as a simple predictor did improve the model fit (improvement to AIC > 2), whereas its inclusion as an interaction term did not (improvement to AIC < 2). The output for the best fitting model is presented in Table 6.10.

	Estimate	Std. Error	*z*-value	*p*-value	
Intercept	5.46	1.01	5.41	<**0.001**	***
Response Target	1.19	0.45	2.66	<**0.01**	**
Predicate	0.47	0.63	0.74	0.46	
Attachment	2.29	1.58	1.45	0.15	
Predicate:Attachment	3.73	1.30	2.87	<**0.01**	**

Table 6.10: Model output for response accuracy.

This indicated a significant effect of *response target* ($\hat{\beta} = 1.19$; $z = 2.66$; $p < 0.01$) with participants giving more accurate responses when the target response was 'no.' No significant effects of *predicate* or

attachment were observed. There was however a significant interaction between the two ($\hat{\beta}$ = 3.73; z = 2.87; $p < 0.01$). To follow up on this significant interaction, we ran pairwise comparisons of *attachment* within *predicate*. This revealed a significant effect of *attachment* within the perceptual items ($\hat{\beta}$ = 4.52; z = 2.19; p = 0.03) which surfaced as greater accuracy for HA items. The effect of *attachment* within the non-perceptual items was non-significant ($\hat{\beta}$ = 0.43; z = 0.29; p = 0.77). Again due to the relatively small number of incorrect responses (N = 35, 4.14%) we inspected, we again inspected which items participants got wrong. This was to explore if the trend above could have been driven by a relatively small number of problematic items. To that end, we calculated the IQR for accuracy by item. For no items, did their accuracy fall 1.5 IQR above/below the third/first quartile.

6.4.3 Interim discussion

Above, we have presented an eye-tracking-while-reading experiment to follow up on the response accuracy results from its sister experiment, the self-paced reading experiment in section 6.3.

Starting with the eye-movement measures in the critical region, results are largely convergent with those in the self-paced reading experiment. Participants read LA-disambiguated non-perceptual items significantly more quickly than HA-disambiguated ones in the FP and TT measures confirming ($H_{6.1}$). For the perceptual items, there was a trend for participants to read HA-disambiguated items more quickly in the TT measure (the measure in which Aguilar et al. 2021 observed their significant PR effect) which was consistent with ($H_{6.2}$). Although this was only marginal (p = 0.06), we should keep in mind that (i) we had noticeably fewer participants in the eye-tracking experiment (27 vs 66) and (ii) the HA bias within the perceptual items was fully significant in both of the analysed windows for the self-paced reading task. Thus we should not take this lack of a fully significant *post hoc* in the eye-tracking study as evidence against an online bias for PRs in the relevant contexts.

Shifting to the post-critical region, it is unclear why the observed LA bias was not restricted to the non-perceptual items as observed in the critical eye-tracking region or the critical and post-critical self-paced reading regions. In isolation, we might take this to suggest an initial locality bias regardless of PR availability. However, we should be cautious of such an interpretation for several reasons. First, our earliest eye-tracking measures indicate that participants are already treating the perceptual and non-perceptual items differently as soon as they encounter the disambiguating region with the locality effect being restricted to true RCs. Second, at both the critical and post-critical self-paced reading windows, the locality effect was restricted to the non-perceptual items. Third, the only other eye-tracking study on this topic, Aguilar et al. (2021), did not report a similar effect in their critical or post-critical regions. As such, the result does not appear to be replicable and we should be wary of trying to interpret the effect.

Turning to the response accuracy, we again observed that after items with perceptual matrix predicates, participants were significantly more accurate in their response to HA items than LA ones. Thus, even though the global error rates in our data were rather small, the fact that the same effect surfaced in both the eye-tracking and self-paced reading experiments despite the differences in methodology and participants, suggests that the pattern is unlikely to be spurious. Moreover, this effect was not attributable to a subset of potentially problematic items in either case. As such, the present study presents a new type of evidence in support of the PR account; not only is a PR-firstness effect observable in the acceptability judgements of fully disambiguated items (Pozniak et al. 2019), interference from PR-firstness is observable even in the interpretation of such items.

The replicability of the persistence effect with PRs leads to a new question: why only with PRs? In the same experiments, we did not observe greater accuracy with LA disambiguated non-perceptual items (resulting in a significant interaction) despite repeated evidence of a LA bias in the online parsing of true RCs. Similarly in chapter 5, we did not find greater accuracy for non-SpecIP disambiguated overt pronominal items,

despite a clear online bias for non-SpecIP antecedents in the same items in both the self-paced reading and eye-tracking-while-reading data. As such, it seems we must say something more than that an intermediate parse may interfere with the response to comprehension questions.

If we focus for just a moment on the PR/RC items, we might speculate that a potential answer to this relates to what is reanalysed. In the case of a HA RC in the non-perceptual items, participants only had to reanalyse where the adjunct RC attaches. In the case of reanalysing a PR to a RC, however, participants may have also needed to reanalyse the argument structure of the matrix CP. Recall from chapter 2 that PRs can occupy the complement of perceptual predicates.[73] Therefore, let us assume that when encountering the embedded tensed element in (85), participants project a PR as the DO. In that case, the DO of the intermediate parse would be a DP denoting an event of running in which the colleague is the agent. However, upon encountering the disambiguating secondary predicate, the PR DO would have to be reanalysed as a DP denoting an individual, in this case, a particular colleague.

(85) a. *Gianni* | *ha visto* | *il collega* | *della biologa* | *che correva* | *sporca* | *di fango.*
Gianni has seen the collegue.M of.the biologist.F that ran dirty.F of mud
'Gianni saw the colleague of the biologist that was running covered in mud.' [Italian]

This potential asymmetry in the reanalysis of the DO is relevant as the literature on persistence effects has focused on sentences such as (86). When the 'the deer' is first encountered in that sentence, it is a potential and plausible DO object for the verb 'hunt.' However, upon encountering the immediately following verb 'ran,' 'the deer' must be reanalysed not as the DO of the adjunct CP, but as the subject of the matrix CP. Despite the global unambiguity of such sentences, previous studies have found that participants still give incorrect responses to the comprehension questions more frequently for these items than control items (Christianson et al. 2001), even in cases where reanalysis is straightforward (Sturt 2007). This suggests that once participants commit to an initial DO parse, that initial semantic representation may persist and interfere with the final interpretation even in cases of syntactic reanalysis. Thus, we might speculate that we observe persistence effects in PR-compatible but not RC-only contexts because only the former requires reanalysis of the initially eventive DO. Nonetheless, as we cannot directly test the 'DO' idea with the present results, we leave identifying the source of persistence effects in PR-compatible contexts to future research.

(86) While the man hunted the deer ran into the woods.
Q. Did the man hunt the deer? Yes or no (Christianson et al. 2001)

The discussion has so far not touched on two asymmetries with the previous work on PR parsing: (i) the lack of adaptation effects in the present studies and (ii) the lack of interference effects with comprehension questions in Lee & De Santo (2022). Let us explore each of these in turn.

Starting with the lack of adaptation effects in the present study, we might suggest that this may relate to the quantity or quality of non-PR/RC items. This is because the present study contains a noticeably higher percentage of other types of items than in Pozniak et al. (2019) (48% vs 69%) and unlike the non-PR/RC items in Aguilar et al. (2021) we included another type of temporarily ambiguous items. Those pronominal items accounted for 46% of our non-critical PR/RC items and contained a different type of ambiguity. Our task was also noticeably shorter than that in Fernandes et al. (2018) (32 vs 60 critical items). However, given that some adaptation effects were observed in the pronominal data, it is not clear that the above-mentioned factors alone can account for the lack of adaptation effects. As such, we leave the issue of adaptation effects in PR parsing open to future work.

[73] We are aware of the fact that Cinque (1992) argues that PRs may sometimes instead adjoin to VP or NP depending on the selectional properties of the matrix predicate. However, given that many of our matrix predicates (e.g., *vedere* - 'see') unambiguously admit VP complement PRs (Cinque 1992) this does not meaningfully affect our argument.

Regarding the fact that the present study observed an offline effect whereas Lee & De Santo (2022) did not, we suggest that this may be due to the manipulations used to force attachment. Whereas Lee & De Santo (2022) exploited a number manipulation, the present paper used gender marking to force attachment. Although we are not aware of any paper that has compared the effect of the two manipulations on attachment in Italian, Slioussar, Antropova & Chernova (2022) recently presented some evidence that, at least in Russian, interference from an initial/preferred parse is more evident with a gender manipulation than a number one. Namely, when they forced participial clauses attachment with or against their HA bias, the authors noted an effect of attachment within the gender-disambiguated items (HA: 89% accuracy; LA 69%). However, the same was not true for items disambiguated by number (HA: 83%; LA 82%). This asymmetry is also in line with the findings in Carminati (2005) for a different type of processing bias. They observed that when null pronouns in Italian are forced to violate their antecedent bias participants were more readily able to overcome the structural preferred antecedent when items were disambiguated by number than when they were disambiguated by gender.

Regardless of the new open questions, the consistency between the results from our first two experiments in this chapter indicates that the self-paced reading task from section 6.3 present a suitable baseline against which to explore attrition effects. To that end, we now turn to a third experiment with native Italian speakers living in a majority English-speaking country.

6.5 Experiment 3: Self-paced reading again

6.5.1 Method

6.5.1.1 Participants

Experimental participants for the self-paced reading task consisted of 31 native speakers of Italian (Female = 24; Male = 7).[74] As above, all participant (i) had lived in Italy from birth until at least the age of 16, (ii) had no diagnosed language-related disorders, and (iii) reported growing up monolingually.[75] At the time of testing, all participants were living in a majority English-speaking country and had been there for a minimum of 5 years (mean = 12.34 years; SD = 7.69 years).[76] Participants were pre-screened to ensure that no participants had travelled back to Italy in the two months prior to testing to avoid any potential re-exposure effects (Chamorro et al. 2015a). At the time of testing the experimental group had a mean age of 39.81 years (SD = 7.69).[77] As with our other experiment, participants were asked to report how many languages they could speak at or above an 'intermediate' level. They were also asked to indicate for how many hours they used each of these languages in a typical day. To quantify participants' language use, we calculated their percentages for each language (i.e., hours of language x / total hours reported for any language). Table 6.11 reports the means and standard deviations for their use of Italian, English, and 'Other' alongside the values for the control group reported in section 6.3.

We also had this group take the Cambridge Assessment General English quick placement test[78] given previous work's focus on 'near-natives' (e.g., Tsimpli et al. 2004). On average the group scored 22.10 / 25

[74] These are slightly different participants from those reported in section 5.5. Due to insufficient data from some participants in the relevant item sets (< 50%), we could not consider all 33 of the original participants for both the pronominal and PR/RC items. As such we do not consider 2 participants for the present analysis as well as 1 participant in the pronominal analysis.

[75] Even though all participants reported growing up monolingually at the pre-screening, when specifically asked about potential exposure to 'dialect' as a child (e.g., with one's grandparents), 8 indicated some exposure.

[76] Note, as all participants reported growing up monolingually in Italy, the experimental participants qualify as first-generation immigrants (i.e., potential attriters), not second-generation immigrants (i.e., heritage speakers).

[77] As mentioned in footnote 42, although our participants are somewhat older than those in Chamorro et al. (2015a) or Martín-Villena (2023), they are more comparable to those in Tsimpli et al. (2004: Ianthi Tsimpli, p.c.). Moreover, in each experiment, our experimental and control groups are comparable in age. As such, age is highly unlikely to drive any potential group-level effects unlike what was reported in Kaltsa et al. (2015) who reported a 30 year age difference between groups.

[78] Items available at: https://www.cambridgeenglish.org/test-your-english/

	Control		Experimental	
	Mean	SD	Mean	SD
Italian	82.85%	14.28%	15.82%	15.79%
English	13.07%	11.84%	82.88%	14.74%
Other	4.08%	6.88%	1.31%	3.73%

Table 6.11: Mean language use by participant as percentages of a typical day with standard deviations for the experimental and control groups.

(SD = 1.97) suggesting they are likely upper-intermediate to advanced L2 speakers of English. Therefore, the experimental group's prolonged residency in their L2 community, lack of recent re-exposure, frequent L2 use, attenuated use of the L1, and high L2 proficiency suggest this is a suitable group in which to look for attrition.

6.5.1.2 Stimuli

The stimuli were identical to those used for the control group. An example critical items (N = 32) in the perceptual LA condition is repeated in (83, partial). As before, the experiment also included the 32 pronominal items reported in chapter 5 as well as 40 unambiguous fillers for a total of 104 sentences.

(83, partial) a. Perceptual LA [Italian]

Gianni | *ha visto* | *il collega* | *della biologa* | *che correva* | *sporca* | *di*
Gianni has seen the colleague.M of.the biologist.F that ran dirty.F of

fango.
mud

'Gianni saw the colleague of the biologist that was running covered in mud.'

6.5.1.3 Procedure

The procedure was identical to the control group reported in section 6.3.1.3 except for the addition of the English placement test. To summarise, participants first provided informed consent and filled out a language history questionnaire. This was followed by the self-paced reading task (window boundaries indicated by '|'s in 83, partial). After each item, participants were presented with a polar comprehension question with the possible responses labelled 'J' and 'F.' Once participants had completed the self-paced reading task, they were then presented with the English placement test.

6.5.1.4 Data cleaning and analysis

As above, we present two planned analyses as well as one unplanned analysis. For the planned analyses, we consider the reading times of the critical window and response accuracy. Given De Santo & Lee's (2022) additional analysis of the post-critical window, we also present an unplanned analysis that considers the reading times of our post-critical window. For these analyses, we combine the data from the Italian living in their home community reported in section 6.3 (henceforth the control group) and the new data from the Italians living in a majority English-speaking country (henceforth the experimental group).

Prior to analysis, we first cleaned the data from the experimental group. To that end, we coded as missing any trial for which the analysed window, or any window prior to it, had an implausibly fast reading time (<

200 ms). This affected 1.21% of trials in both the critical and post-critical windows. For the two windows, we then calculated the IQR per condition and coded as missing any value that lay 1.5 IQR above the third quartile for that window. This affected a further 7.86 % of trials in the critical window (for a total data loss of 9.07%) and 9.78% of trials in the post-critical window (total data loss: 10.99%). The remaining data was equally distributed across the 4 conditions (critical: $\chi^2(1,902) = 0.18$; $p = 0.98$; post-critical: $\chi^2(1,883) = 0.26$; $p = 0.97$). For the analysis of the comprehension questions, we only consider those items included in the planned analysis of the critical window.

To analyse the data, we used the *lme4* package (Bates et al. 2015) in R (R Core Team 2022). We followed the procedure described above to identify the best random-effects structure. Namely, we conducted families of intercept-only models in which we considered the various random-effects structures of our theoretically relevant fixed predictors and selected the best fitting random effects structure (Matuschek et al. 2017) using the AIC. We then fit a base model with only the fixed effects that we were theoretically interested in using sum coding (i.e., -0.5, 0.5). For *attachment*, we coded LA as the negative level. For *predicate*, we coded non-perceptual as the negative level. For *group*, we coded control as the negative level. To account for potential adaptation effects, we then conducted two further models which included *order* (centred over the experiment) as either a simple predictor or potential interaction term. If complicating the model reduced the AIC by 2 or more, the more complicated fixed-effect structure was maintained. For the analysis of the comprehension question responses, we then repeated this step with *response target* (negative level: yes) to account for any potential effects of the acquiescence bias (Holbrook 2008).

6.5.1.5 Hypotheses and predictions

The following hypotheses were drawn:

($H_{6.3}$) Overall, participants will exhibit an online sensitivity to PR availability that is not affected by attrition.

($H_{6.4}$) Globally, the experimental group will exhibit a shift toward a LA bias when compared to the control group.

For reading times (in both windows) as well as response accuracy, we expect a significant interaction of *attachment* and *predicate*. This is predicted to be driven by a LA bias for items that unambiguously contain RCs (Late Closure) as well as a HA bias for items that are locally compatible with PRs (PR-First Hypothesis). Given ($H_{6.3}$), this is expected to surface regardless of any effect of *group*. This was empirically motivated by the results reported in chapter 3. Recall that in that experiment, both groups exhibited a clear effect of PR availability with a HA bias in PR-compatible contexts and a LA bias in PR-incompatible contexts. This was despite a global shift toward LA in the experimental group. Integrating ($H_{6.4}$), we predict a two-way interaction of *attachment* and *group* abstracting away from PR-availability. This is expected to surface as a shift toward a LA bias in the experimental group when compared to the experimental group. Note, given our discussion in chapter 3 about how the perceptual items are still compatible with RC parses, we do not go so far as to expect a three-way interaction of *attachment*, *predicate*, and *group*.

6.5.2 Results

6.5.2.1 Reading times in the critical window

Figure 6.7 presents the mean reading times of the critical window for the experimental group alongside the control group for reference. On average the experimental group read HA non-perceptual items (mean = 873.71 ms) slower than LA ones (mean = 718.01 ms). For perceptual items reading times for the two attachments were similar (HA: mean = 767.68 ms; LA: mean = 786.93).

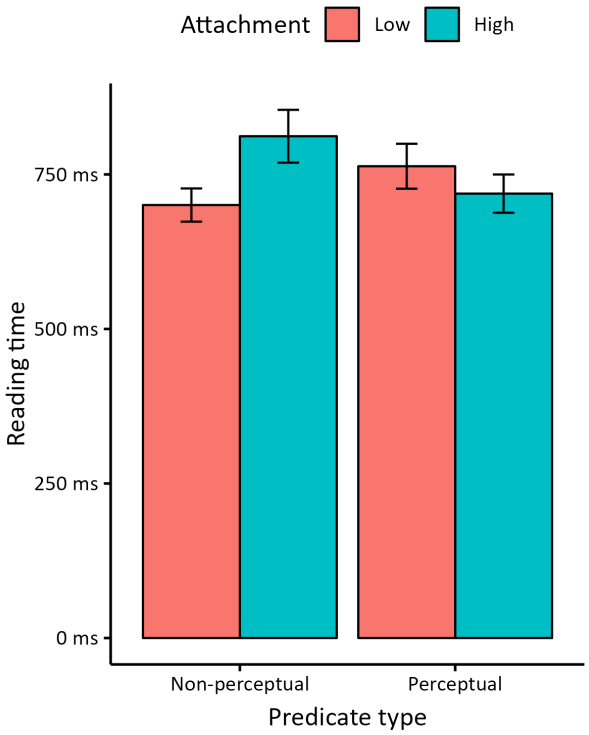

(a) Control

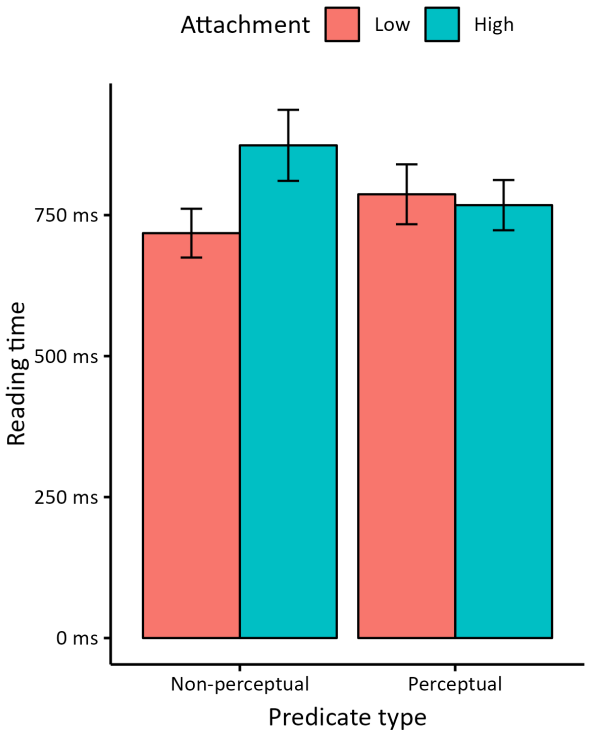

(b) Experimental

Figure 6.7: Global average reading times for the critical window for the control and experimental groups with 95% confidence intervals.

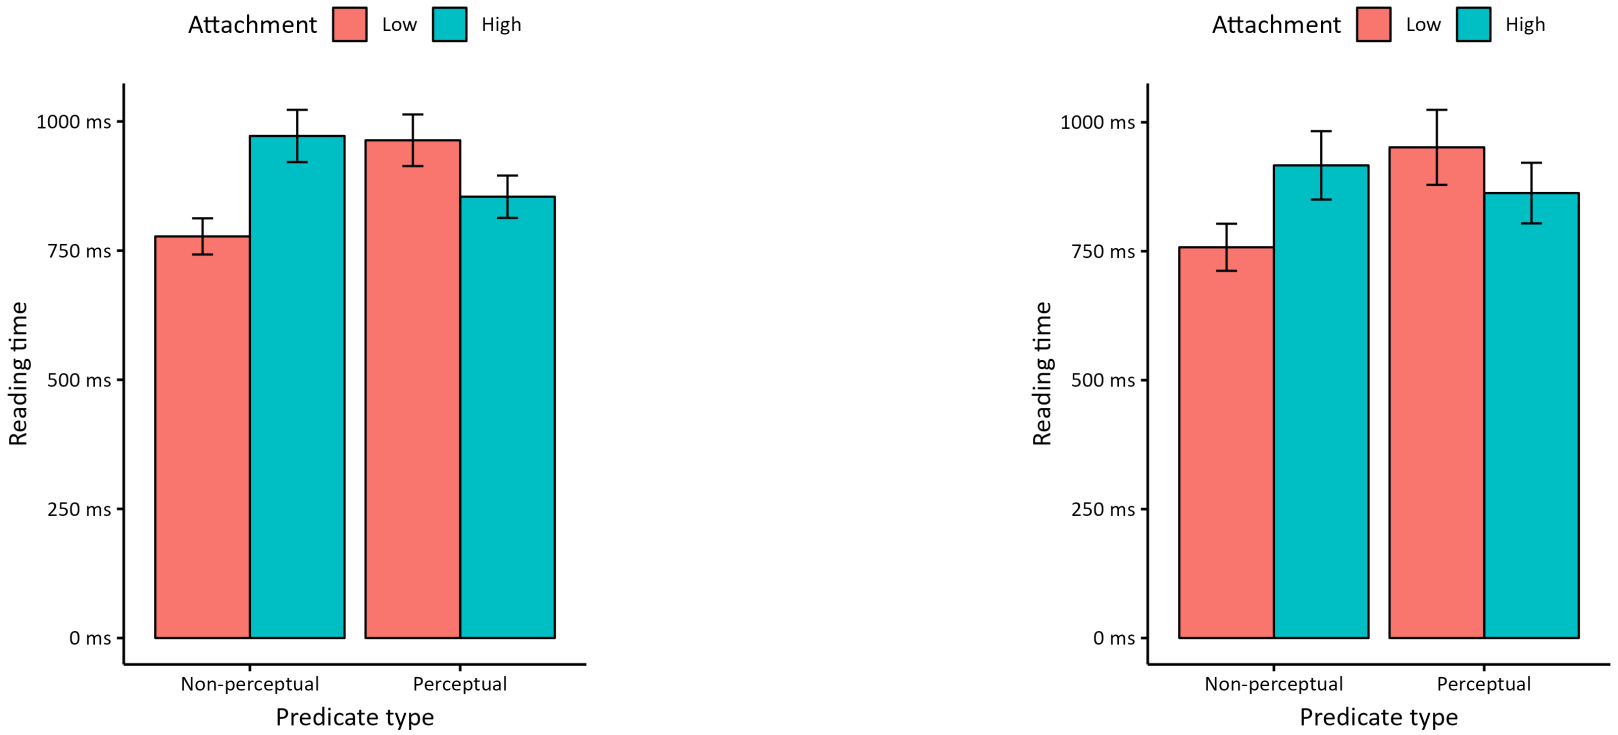

(a) Control

(b) Experimental

Figure 6.8: Global average reading times for the post-critical window for the control and experimental groups with 95% confidence intervals.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.55	0.03	189.51	$<$ **0.001**	***
Predicate	-0.01	0.01	-0.92	0.36	
Attachment	0.05	0.01	3.31	$<$ **0.001**	***
Group	0.06	0.07	0.95	0.34	
Order	-0.001	0.0002	-6.16	$<$ **0.001**	***
Predicate:Attachment	-0.16	0.03	5.66	$<$ **0.001**	***
Predicate:Group	0.002	0.03	0.07	0.95	
Attachment:Group	0.04	0.03	1.39	0.17	
Predicate:Attachment:Group	-0.03	0.06	-0.49	0.62	

Table 6.12: Model output for the reading times of the critical window for the experimental and control groups.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.67	0.03	191.37	$<$ **0.001**	***
Predicate	0.04	0.02	2.70	$<$ **0.01**	**
Attachment	0.04	0.02	2.71	$<$ **0.01**	**
Group	0.003	0.07	0.43	0.97	
Order	-0.003	0.0003	-10.21	$<$ **0.001**	***
Predicate:Attachment	-0.26	0.03	-8.13	$<$ **0.001**	***
Predicate:Group	0.04	0.03	1.35	0.18	
Attachment:Group	-0.01	0.03	-0.35	0.73	
Predicate:Attachment:Group	0.07	0.07	1.08	0.28	

Table 6.13: Model output for the reading times of the post-critical window for the experimental and control groups.

Reading times for both groups were logged and subject to a family of mixed-effect regression as discussed above. Model comparison indicated that the best fitting model included *order* as a simple effect but not as a potential interaction term. The output for that model is reported in Table 6.12. That model indicated a significant effect of *attachment* ($\hat{\beta}$ = 0.05; t = 3.31; $p < 0.001$) indicating that participants read HA-disambiguated items more slowly. There was also a significant effect of *order* ($\hat{\beta}$ = -0.001; t = -6.16; $p <$ 0.001) with participants reading more quickly as the experiment progressed. The effects of *predicate* and *group* were non-significant. However there was a significant two-way interaction of *predicate* and *attachment* ($\hat{\beta}$ = -0.16; t = 5.66; $p < 0.001$). To follow up on that interaction we conducted pairwise comparisons of *attachment* within *predicate* with Holm-Bonferroni corrected *p*-values. That indicated a significant effect of *attachment* within the non-perceptual items ($\hat{\beta}$ = 0.12; t = 6.34; $p < 0.001$) with participant reading LA-disambiguated items more quickly. The effect of *attachment* was non-significant within the perceptual items ($\hat{\beta}$ = -0.03; t = -1.67; p = 0.010). None of the interaction terms with *group* were significant.

6.5.2.2 Reading times in the post-critical window

Figure 6.8 presents the mean reading times of the post-critical window for the experimental group alongside the control group for reference. On average the experimental group read HA non-perceptual items (mean = 916.48 ms) slower than LA ones (mean = 757.67 ms). For perceptual items reading times HA-disambiguated items (mean = 862.77 ms) were instead read more quickly than LA-disambiguated ones (951.43 ms).

Reading times for both groups were logged and subject to a family of mixed-effect regressions. Model comparison indicated that the best fitting model contained *order* as a simple effect. The output for that model is reported in Table 6.13. The model indicated a significant effect of *predicate* ($\hat{\beta}$ = 0.04; t = 2.70; p < 0.001) indicating that participants read perceptual items more slowly. There was also a significant effect of *attachment* ($\hat{\beta}$ = 04; t = 2.71; $p < 0.001$) indicating that participants read HA-disambiguated items more slowly. The model also indicated a significant effect of *order* ($\hat{\beta}$ = -0.003; t = -10.21; $p < 0.001$) with participants reading more quickly as the experiment progressed. The effect of *group* was non-significant, as were all of the interaction terms that included it. The only significant interaction was between *attachment* and *predicate* ($\hat{\beta}$ = -0.26; t = -8.13; $p < 0.001$). To follow up on that interaction we conducted pairwise comparisons of *attachment* within *predicate*. For these, we report Holm-Bonferroni corrected *p*-values. That indicated a significant effect of *attachment* within the non-perceptual items ($\hat{\beta}$ = 0.18; t = 7.60; $p < 0.001$) with faster reading times for LA-disambiguated items. The effect of *attachment* was also significant within the perceptual items ($\hat{\beta}$ = -0.09; t = -3.87; $p < 0.001$) indicating that participants read LA-disambiguated items more slowly.

6.5.2.3 Response accuracy for comprehension questions

Figure 6.9 presents the average response accuracy for the experimental group alongside the control group for reference. In the non-perceptual items, experimental participants gave the correct response at similar rates for LA (mean = 91.44%) and HA (89.38%) disambiguated items. For the perceptual items, responses were noticeably more accurate for HA items (96.52%) than LA items (83.24%).

Responses (coded as ± correct) were subjected to a family of mixed-effect logistic regressions as described above. Model comparison indicated that the inclusion of *order* did not improve the mode, however, the inclusion of *response target* as a simple predictor did. That model is presented in Table 6.14. The model indicated a significant effect of *response target* ($\hat{\beta}$ = -0.41; z = -2.89; $p < 0.01$) with participants more accurate when the response target was *sì*. The effects of *predicate*, *attachment*, and *group* were non-significant. The model also indicated that all of the potential interactions with *group* were non-significant. The interaction of *predicate* by *attachment* was significant however, ($\hat{\beta}$ = 2.15; z = 6.47; $p < 0.001$).

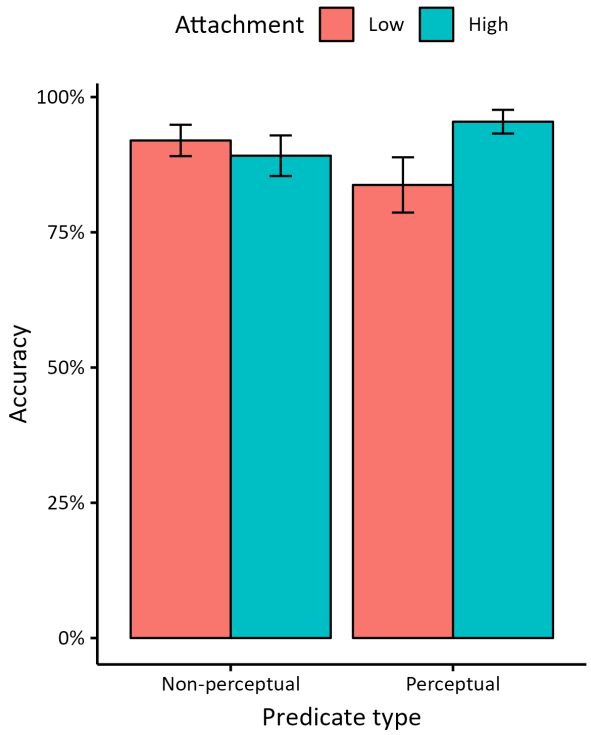

(a) Control

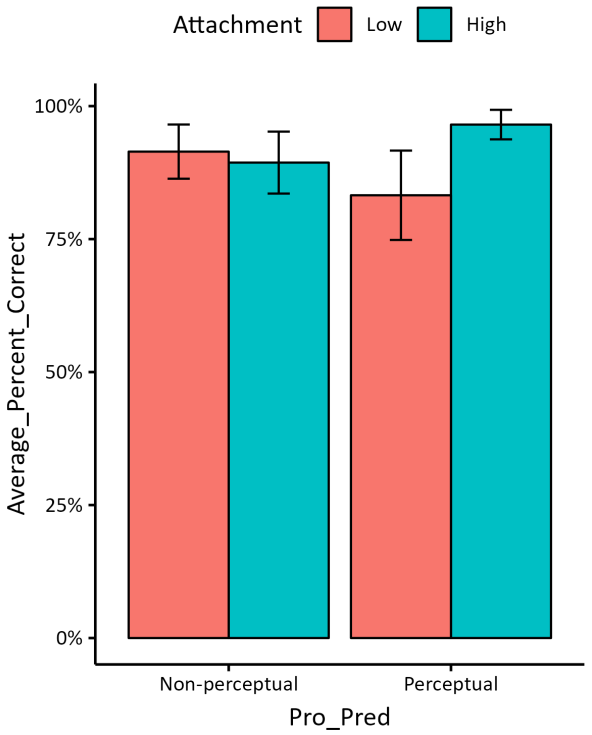

(b) Experimental

Figure 6.9: Average response accuracy (by participant) for the control and experimental groups with 95% confidence intervals.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	2.98	0.02	15.94	$< \mathbf{0.001}$	***
Predicate	0.15	0.17	0.98	0.37	
Attachment	0.57	0.32	1.81	0.07	.
Group	0.14	0.33	0.44	0.66	
Response target	-0.41	0.14	-2.89	$< \mathbf{0.01}$	**
Predicate:Attachment	2.15	0.33	6.47	$< \mathbf{0.001}$	***
Predicate:Group	0.21	0.33	0.65	0.52	
Attachment:Group	0.19	0.40	0.48	0.65	
Predicate:Attachment:Group	0.31	0.66	0.48	0.63	

Table 6.14: Model output for the response accuracy for the experimental and control groups.

To follow up on that interaction we conducted pairwise comparisons of *attachment* within *predicate* using the *emmeans* package (Lenth 2022). For this, we report Holm-Bonferroni corrected *p*-values. Those comparisons indicate that the effect of *attachment* within the non-perceptual items was non-significant ($\hat{\beta}$ = -0.51; z = -1.48; $p = 0.14$). The effect of *attachment* within the perceptual items however was significant ($\hat{\beta}$ = 1.65; z = 4.44; $p < 0.001$) with participants giving more correct responses with HA-disambiguated items. As in section 6.3 we inspected accuracy by item to see if the trend could have been driven by a problematic set of items. To that end, we calculated the IQR we calculated the IQR for accuracy by item. For no items, did accuracy fall 1.5 IQR below/above the 1st/3rd quartile.

6.6 Discussion

Above we have presented an experimental group of Italians living in a majority English-speaking country for our self-paced reading task. Results, from these participants were then compared against those of the control group reported in section 6.3 to investigate how the online resolution of PR/RC parsing ambiguities are affected by attrition. Specifically, we are interested in whether the experimental group exhibits a global shift toward a LA bias (as compared to the control group, $H_{6.4}$) despite a continued online sensitivity to PR-availability ($H_{6.3}$).

Overall results indicated that our experimental group resolved PR/RC parsing ambiguities in a similar way to the control group. In no measure, did we observe a significant effect of, or interaction with, *group*. Nor did the results from the pooled data meaningfully differ from those for the control group when they were analysed by themselves. Starting with the comprehension responses, participants exhibited similar levels of accuracy for non-perceptual items regardless of attachment. For perceptual items, however, there was a clear trend for greater accuracy after HA disambiguated items, which we have interpreted as a persistence effect with PRs. Turning to reading times in the post-critical window, we again observed a clear LA bias for non-perceptual items and a HA bias for perceptual ones.

The single caveat to the parallels between the pooled data and the control-only data arises from the critical window reading times. There, we did not observe a global HA bias for perceptual items, despite a clear LA bias with non-perceptual ones. This raises the question of how we should interpret this lack of a HA bias. On the one hand, we might interpret this to mean that the experimental group exhibited a weaker PR bias in the critical window. Due to a lack of a significant interaction with *group*, however, we discard this interpretation. On the other hand, we might suggest that the significant HA bias for the control group in the critical window

was spurious. However, this position seems unlikely given the significant PR effects in the post-critical window and comprehension responses, both of which are replicated in the pooled data. As another alternative, it may be that the inclusion of the experimental group also introduced additional variance. After all, even though both groups consist of native Italian speakers, there are striking differences in their actual input; in a typical day, the control group uses Italian substantially more frequently than the experimental group (control mean = 82.85%; experimental mean = 15.82%; see Table 6.11). This combined with the fact that the effect for the control group was already relatively small at the critical window (LA mean = 763.24 ms; HA mean = 718.97 ms; difference 44.27 ms) may have obfuscated the pattern we are interested in. As for why the inclusion of the experimental group did not wash out the PR effect in the post-critical window, a natural response assuming this interpretation would be that the effect in that window was already noticeably larger for the control group (LA mean = 963.56 ms; HA mean = 854.33 ms; difference 109.23 ms).

Regardless, the lack of any significant interaction means that, *contra* ($H_{6.4}$), we failed to replicate the global shift toward more LA reported by Dussias (2003, 2004) and Dussias & Sagarra (2007) for Spanish speakers immersed in English. An immediate question we might raise is whether this is somehow attributable to our participant sample. This seems unlikely for several reasons. First, as noted in section 6.5.1.1 the participants' prolonged L2 immersion, lack of recent re-exposure to the L1, frequent L2 use, attenuated L1 use, and high L2 proficiency all suggest this is an ideal group in which to explore attrition effects. Second, although it is not straightforward to directly compare our background measures to those in Dussias (2003, et seq.) as noted in chapter 3 we can compare the length of residency in the L2. Doing so we observe that our participants had been immersed for longer (e.g., Dussias 2003 mean = 7.5 years; current sample mean = 12.34 years). This is relevant as Dussias & Sagarra (2007) found that, for their sample, length of L2 immersion was a significant predictor of attrition within advanced L2 speakers. Third, we can directly compare our present experimental group with the one reported in the sentence interpretation task with globally ambiguous items in chapter 3. Recall that there we did observe a global shift toward LA which we interpreted as an increased locality effect for true RCs. Between the experimental groups, language use in a typical day was similar (chapter 3 English mean = 76.16%; current sample English mean = 82.88%). As was their L2-English proficiency (chapter 3 mean = 21.76 / 25; current sample mean = 22.10 / 25) and length of residency (chapter 3 mean = 14.27 years; current sample mean = 12.34 years). Fourth, there was non-trivial overlap in the experimental participants presented in this chapter and those in chapter 5; of the present 31 experimental participants, data from 30 was also considered in the pronominal analysis (out of a total of 32). There again we observed an attrition effect in the processing of overt pronouns.

If we cannot attribute the lack of an effect to our participants, we might want to consider the items. Recall that in chapter 3, we suggested that our items may have (partially) contributed to the difference in the size of the shift toward LA interpretations between Dussias (2003) and our own results. This was because the attrition effect appears to have surfaced as an increased locality effect for true RCs. For our unambiguously RC items, the control group in chapter 3 selected HA interpretation on average 30.52% of trials. As such, this likely imposed a logical limit on the size of the possibly observable attrition effect given that attrition did not lead to a collapse of the PR/RC distinction (i.e., the HA bias was stable in the PR-compatible contexts) and the rate of HA could not be negative. However, we cannot apply a similar line of reasoning to explain the lack of a group effect in the online measures given the (functionally) continuous nature of reading times. As such, it remains unclear why we were unable to detect an attrition effect. However, we should keep in mind that the lack of an effect in the present experiment should not be taken as evidence that these biases do not attrite in general given the previous results from the sentence interpretation tasks in Dussias (2003) and chapter 3 as well as the online processing results from Dussias (2003, 2004) and Dussias & Sagarra (2007).

As a final note on ($H_{6.3}$), although we observed a global LA bias within non-perceptual items that clearly differed from the bias within perceptual items and although these biases were not affected by group, we cannot unambiguously conclude that PR-sensitivity is unaffected by attrition. This falls out trivially from the lack of

identifiable attrition effects in the experimental group. It is still possible that the HA to LA shift observed in Dussias (2003, et seq.) may have been driven by a change in the parsing of PRs.

6.7 Chapter summary

In this chapter, we have presented three experiments with locally ambiguous PR/RC items. As there was relatively little prior work on the online parsing of PR/RC ambiguities, the first two experiments aimed to investigate some predictions drawn from the PR-First Hypothesis. Results from those experiments largely confirm its predictions for Italian. Despite being classified as a 'HA language,' once PR availability is controlled for, RC in Italian exhibited the expected LA bias. A HA bias was only observed in PR-compatible contexts. Moreover, our results also provided a novel form of evidence in support of the PR-First Hypothesis. Namely, in contexts that are temporarily PR-compatible, even once the sentence is fully disambiguated away from a PR, there is a dip in response accuracy which we have interpreted as a persistence of the intermediate PR parse. Against that backdrop, our third experiment was intended to explore how attrition affected PR/RC parsing biases. However, despite our background measures indicating our experimental group was an ideal group in which to look at attrition, no changes to PR/RC parsing biases were observed.

7 General discussion and conclusions

7.1 Introduction

The present chapter summarises the findings of the thesis and discusses their implications. To that end, section 7.2 first reiterates the overall aims of the thesis as well as the core results from our experiments focusing on attrition. Section 7.3 then integrates these findings with our overarching research question (RQ_0). That section also discusses our results' implication for the source of attrition, directions for future research, and potential points of interest for other areas of linguistic research, before section 7.4 presents our conclusions.

7.2 Summary

7.2.1 Background and aims

In chapter 2, we presented a review of the previous work on attrition in null-subject languages and observed two general patterns. On the one hand, various researchers have reported an apparent relaxation of discourse conditions on the syntactic options offered by null-subject languages (e.g., the [+TS] specification for overt pronoun). On the other hand, we argued that the results from many of those same studies also provided evidence that speakers may additionally exhibit an apparent strengthening of other biases already present in their L1 (e.g., the [−TS] specification for null pronoun). We then pointed out that this second pattern cannot be ascribed to transfer from the L2 in the cases considered, nor could it be accounted for in terms of overt pronouns as a 'processing default.' Consequently, we suggested that attrition may lead to a general increased reliance on one's interpretive biases. As attrition under this idea would not be tied to the conflict between discourse feature specification in a speaker's L1 and L2 nor would it be tied to the overextension of overt pronouns as a processing default under cognitive load, this led us to consider whether similar patterns could be observed beyond the previously studied interface phenomena. Specifically, we focused on 'RC' attachment ambiguities. We chose this phenomenon for two reasons. First, we argued that the interpretation of such ambiguities is conditioned by principles of computational efficiency at the level of the parser rather than discourse features, allowing us to disentangle the role of biases from the syntax-discourse interface. Second, there was some initial evidence from Spanish that these biases might be affected. However, as those earlier experiments predated Grillo (2012), they did not consider the potential confound posed by PRs. Therefore, it was unclear how to interpret their results and integrate them with work on the attrition of interface phenomena without further study. Against that background, we proposed the starting, general research question in (RQ_0).

(RQ_0) Does attrition affect syntax-discourse structures and attachment ambiguities similarly?

If there is something about more external interface phenomena causing them to be susceptible to attrite (e.g., in the case of null and overt pronouns, that they require the integration of syntactic and discourse information), then we would expect pronominal reference and attachment ambiguities to be affected differently. However, if attrition can lead to an increased reliance on one's L1 biases more generally, then we would expect similar attrition effects in the parsing of attachment ambiguities despite the difference in interface status. In order to tackle this question, we conducted a series of experiments to which we now turn.

7.2.2 Experimental results

7.2.2.1 Interpretive biases

In chapter 3, we presented a sentence interpretation task to compare how attrition affects native Italian speakers' interpretation of pronominal reference and PR/RC parsing ambiguities. For this task, we employed globally-ambiguous items already present in the literature. Pronominal items were minimally adapted from Tsimpli et al. (2004) and varied by direction (matrix- or embedded-first) as well as pronominal form (null or overt). PR/RC items were taken from Grillo & Costa (2014). Within those items, only the matrix predicate (perceptual or non-perceptual) was manipulated. Participants consisted of two groups of native Italian speakers: a control group still living in Italy and an experimental group living in a majority English-speaking country. For the second group, participants were pre-screened and asked to fill out a language background questionnaire to ensure they were a suitable group in which to look for attrition effects. After reading each item in the sentence interpretation task, participants were asked to respond to comprehension questions by selecting their preferred interpretation of the ambiguous element.

Pronominal results indicated the experimental group did not differ from the control group in their interpretation of embedded-first items. For items with null pronouns, both groups exhibited an overwhelming preference to interpret the pronominal subject as coreferential with the potential referent in matrix SpecIP. For items with overt pronouns, both groups instead exhibited a non-SpecIP bias.

For the matrix-first items, the two groups again patterned together in their interpretation of overt pronouns, both exhibiting a clear non-SpecIP bias. However, the two groups were found to differ in their interpretation of matrix-first items with null pronouns. There, the control group selected SpecIP interpretations at around chance levels, whereas the experimental group selected SpecIP interpretations noticeably more frequently resulting in a significant *group* by *pronoun* interaction.

Results from the PR/RC items indicated that the control group was sensitive to the availability of PRs. When PRs were locally available (i.e., in the perceptual items), that group exhibited a HA bias. When PRs were locally blocked by the selection properties of the matrix predicate (i.e., in the non-perceptual items), however, the same group exhibited a LA bias. A similar pattern was observed in the experimental group's interpretations. Nonetheless, that group was found to select HA interpretations at a significantly lower rate overall.

7.2.2.2 Processing biases

In chapter 5 and 6, we presented a series of 3 experiments to compare how attrition affects the online processing of our two phenomena. For this, we used the novel pronominal and PR/RC items developed in chapter 4. Within the pronominal items, we manipulated the pronominal form (null or overt) and forced its antecedent (SpecIP or non-SpecIP) via gender marking. Within the PR/RC items, we manipulated PR availability via the matrix predicate (perceptual or non-perceptual) and forced the attachment of the embedded CP (HA or LA) again via gender marking.

In Experiment 1, items were presented as part of a self-paced reading task. After reading each item, participants then responded to a polar comprehension question targeting the resolution of the local ambiguity in order to detect inference effects from their initial/preferred parse. For this experiment, participants consisted of native Italian speakers still living in Italy only. This was because we conducted Experiment 1 to first (i) establish a monolingual baseline for our novel temporarily-ambiguous items and (ii) investigate under-researched predictions for monolingual processing drawn from the PR-First Hypothesis. Due to some unexpected results from Experiment 1, we additionally presented a sister experiment (Experiment 2) using an eye-tracking-while-reading task. This was carried out as a validation of the initial experiment (albeit with a different methodology). To that end, Experiment 2 employed the same items and polar comprehension

questions, but the participants consisted of a distinct sample of native Italian speakers still living in Italy. In Experiment 3, we conducted the same self-paced reading task from Experiment 1. For that experiment, however, participants consisted of native Italian speakers living in a majority English-speaking country. This was done to investigate how the online processing or pronominal reference and PR/RC parsing ambiguities are affected by attrition. As in the sentence interpretation task, we prescreened experimental participants and used a language background questionnaire to ensure that they were a suitable group in which to explore attrition. Results from Experiment 3 were then directly compared against those from Experiment 1.

Pronominal results from Experiment 1 indicated that the native Italians still living in Italy exhibited the expected PAH bias for overt pronouns. For those items, they read the disambiguating window significantly more slowly when the pronoun was forced to corefer with the potential antecedent in matrix SpecIP than when it was forced to corefer with a structurally lower potential antecedent. For null pronouns, however, we did not observe the SpecIP bias expected by the PAH. Instead, participants' reading times for SpecIP and non-SpecIP disambiguated items did not differ statistically. We also did not observe any significant effects in the accuracy of participants' responses to the comprehension questions as their accuracy was at ceiling in all conditions.

Results from Experiment 2 corroborated those from Experiment 1. The second group of native Italian participants still living in Italy displayed a non-SpecIP bias for overt pronouns as expected from the PAH. However, we did not observe the expected SpecIP bias for null pronouns. In no eye-tracking measure did the reading times for the critical interest area for those items differ statistically between SpecIP and non-SpecIP disambiguations. Accuracy to comprehension questions was again at ceiling for all conditions and no statistically significant effects were observed.

When the results from Experiment 3 were modelled together with those from Experiment 1 to explore potential attrition effects, this revealed that overall participants exhibited the expected bias for overt pronouns similar to what was observed in Experiments 1 and 2. That is to say, participants read non-SpecIP disambiguated items significantly faster than SpecIP disambiguated items. However, this effect was modulated by two-way interaction with (trial) *order* and a three-way interaction with *order* and *group*. The two-way interaction indicated that although participants generally read more quickly as the experiment progressed (resulting in a significant main effect of *order*) this effect was more pronounced for SpecIP disambiguated items such that the participants were more sensitive to the effect of *antecedent* (i.e., SpecIP or non-SpecIP) toward the beginning of the experiment than they were toward the end. Integrating the three-way effect, the differential effect of *order* on *antecedent* was found to be larger in the experimental group than in the control group. That is to say, the experimental group exhibited a larger non-SpecIP bias for overt pronouns than the control group toward the beginning of the experiment with both groups exhibiting reduced sensitivity to *antecedent* toward the end of the experiment. Thus, although the expected two-way interaction of *group* with *antecedent* was not significant in the final model, we nonetheless observed evidence of attrition which again surfaced as the strengthening of an L1 bias. Furthermore, it is worth underscoring that our significant attrition effect was not dependent on the inclusion of *order* in our statistical modelling. When we additionally conducted a simplified model without *order*, the expected two-way interaction of *group* by *antecedent* was fully significant. Finally, accuracy to the comprehension questions was again found to be at ceiling in all four conditions with no significant difference between the groups.

In Experiment 1, results from the PR/RC items indicated a significant interaction between *predicate* and *attachment* in the critical window. This was found to be driven by opposing biases depending on the local availability of PRs. In the non-perceptual items in which the embedded CP may only be interpreted as a RC, participants exhibited an effect of locality. That is to say, the native Italian speakers still living in Italy were found to read HA-disambiguated items significantly slower than LA-disambiguated ones. In the perceptual items (i.e., PR-compatible items), however, we observed the opposite pattern. Participants read HA-disambiguated items more quickly than LA-disambiguated ones. A similar but more pronounced pattern

was also observed at the post-critical window. In that same experiment, we also observed a significant interaction between *predicate* and *attachment* in the accuracy of participants' responses. This was found to be driven by the items with perceptual predicates. There, responses were significantly more accurate for HA-disambiguated items than LA-disambiguated ones.

Experiment 2 revealed similar results. In the critical interest area, we observed significant interactions of *attachment* by *predicate*. These were driven by a significant LA bias in the non-perceptual items. No such interaction was found in the post-critical interest area where we instead observed a global LA bias in early processing. In the responses to the comprehension questions, however, we again observed a significant interaction between *predicate* and *attachment* driven by a HA bias in the perceptual items. There, participants were significantly more accurate in their responses after HA-disambiguated items than LA-disambiguated ones.

When the self-paced reading task was re-run with native Italian speakers living abroad and compared against the results from Experiment 1, this revealed a significant two-way interaction of *predicate* by *attachment* at the critical region which was driven by a significant LA bias in the non-perceptual items. The HA bias within the perceptual items was non-significant. At the post-critical window, we again observed a significant two-way interaction of *predicate* by *attachment*. However, in that case, the interaction was driven by both a significant LA bias in the non-perceptual items and a HA bias in the perceptual ones. In neither window did we observe an effect or interaction of *group*. As for response accuracy, the model revealed no significant differences between the control group and the experimental group. The model only indicated a two-way interaction of *predicate* by *attachment* driven by greater accuracy for HA-disambiguated items than LA-disambiguated ones in the perceptual items.

7.3 Discussion

Circling back to our starting research question (RQ_0), we take the results from the present thesis to indicate that our response should be positive. This is because attrition effects were observed to affect the resolution of both the pronominal and PR/RC ambiguities. Moreover, whenever such effects were observed, they surfaced as an increased reliance on interpretive/processing biases already present in the L1. That is to say, both phenomena were affected in a qualitatively similar way. The clearest argument for this can be drawn from our sentence interpretation task where we observed attrition effects for both phenomena within the same participants.

Starting with the PR/RC items, one might interpret the significant overall shift toward more LA responses in the experimental participants to indicate a weakening of a 'global' HA bias (cf. Dussias 2003). However, as pointed out above, it is clear that this cannot be the case once PR availability is taken into account. When PRs were locally blocked, participants exhibited a clear LA bias. When PRs were locally available, participants instead exhibited a HA bias. These results would be unexpected if Italian exhibited a global HA bias, but are directly predicted by the PR Account. Thus, we take this as evidence that underlying Italian exhibits two biases: a LA bias for true RCs similar to a 'LA language' like English as well as a bias for PRs over RCs (PRFH), resulting in an apparent HA bias for RCs. As such, one might try to recast the weakening of a 'global' HA bias as a weakening of the bias for PRs as we hypothesised in chapter 3 ($H_{3.5}$). However, based on our results, this too cannot be the case. If attrition were driven by a reduced bias for PRs, we would have expected this to be restricted to our perceptual items, resulting in a significant *group* by *predicate* interaction. This was not the case. Moreover, the experimental group continued to exhibit a clear HA bias in PR-compatible contexts despite a clear LA bias in PR-incompatible contexts and despite a global shift toward more LA interpretations. Therefore, we suggested that the global shift toward more LA responses should be interpreted as an increased reliance on the LA bias for true RCs. This is because, within our experiment, any string that was compatible with a PR parse was also trivially compatible with a RC one. To make this point,

we highlighted the fact that, even within the control group, PR-compatible strings received LA responses in a non-negligible portion of trials. As LA was incompatible with a PR reading in our items, we took this as evidence that despite the PRFH, at least 27.95% of our PR-compatible items were parsed as containing RCs. Moreover, we argued that the true rate of RC parses in PR-compatible contexts was likely higher. This is because unambiguously RC-only strings did not categorically receive LA responses. Therefore, under the reasonable assumption that participants sometimes parsed perceptual items as containing a HA RC, the PR-compatible strings would have also been susceptible to any increased LA bias for true RCs, driving the observed across-the-board effect.

Integrating the pronominal items from the same experiment, we discussed two potential interpretations of the significant *pronoun* by *group* interaction observed in the interpretations of the matrix-first items. On the one hand, we might interpret that interaction to be driven exclusively by an increased SpecIP bias for null pronouns in the experimental group. On the other hand, we might interpret the interaction to be driven by the numerically more exaggerated biases for both null and overt pronouns, despite the trend being noticeably smaller in the latter. Regardless of which interpretation the reader takes to be more likely, our argument remains unchanged. At least for null pronouns, the experimental participants relied more heavily on their interpretive bias than the control group did consistent with ($H_{3.2}$). Although this result may at first appear surprising given that previous work has generally focused on overt pronouns, as argued extensively in chapter 2, it is not without precedent. In fact, our results for null pronouns replicate those reported in Tsimpli et al. (2004) very closely. Moreover, if the reader is willing to accept our re-interpretation of the Gürel's (2004) and Gürel & Yılmaz's (2011) results presented in chapter 2, there is also evidence of a related, replicable pattern in Turkish speakers undergoing attrition. As such, this effect should not be taken as spurious.

Thus, for both of the phenomena investigated, we observe an apparent strengthening of an L1 bias. This in and of itself is noteworthy as it indicates that there is nothing special about the syntax-discourse interface. If structures at more external interfaces such as the syntax-discourse interface were more susceptible to attrition as suggested by the IH, then we would have expected our attrition effect to have been restricted to, or at least noticeably more pronounced in, our pronominal items. Results from our sentence-interpretation task do not provide evidence in support of that position. Rather, our results are more consistent with the idea that any structure which is interpreted non-categorically but is at the same time subject to an interpretive bias may be subject to attrition effects, regardless of that structure's interface status. Of course, this is a very strong position to take and we may need to restrict it as researchers investigate more structures beyond the traditional interface phenomena. Nonetheless, this is a useful starting point for our discussion.

As for the source of our observed attrition effects, in neither of the phenomena we considered, could the effect be reasonably attributed to transfer (either at the level of processing or representations) from the L2. In the case of pronouns, the participants' L2 (i.e., English) lacks an equivalent null form, and as such, there is nothing to be transferred. In the case of the LA bias for RCs, although the L2 does have RCs which also display a LA bias, we discussed the fact that it would be theoretically undesirable to suggest that languages differ in their LA biases for RCs. Moreover, we also discussed the fact that when we consider the monolingual literature through the lens of the PR Account, we have no evidence to believe that the LA bias for RCs in English is stronger than the LA bias in Italian. In fact, when Grillo et al. (2015) conducted the same sentence interpretation task with translational equivalents of the Italian items in Grillo & Costa (2014), the English-speaking participants were not found to differ significantly from the Italian-speaking participant in their interpretation of the non-perceptual (i.e., RC-only) items. That was despite a clear difference in the participants' interpretations of perceptual items which were compatible with PRs in Italian.

Having ruled out the possibility that the increased reliance on interpretive biases was driven by transfer, it appears that we should interpret our group-level differences to be indicative of a 'bilingualism effect.' This is worth highlighting because it suggests that the qualitative similarity in the way that attrition affected the LA bias for RCs and the SpecIP bias for null pronouns is unlikely to have been accidental. In turn, this suggests

that we should seek a parsimonious account under which these two effects receive a unified explanation.

Nonetheless, it is also clear that our results cannot be accounted for by the previously discussed processing-couched accounts such as Sorace's (2011, 2016) idea that overt pronouns function as a 'processing default'. Recall that Sorace (2011) suggested bilingual speakers of Italian's processing of pronominal forms is not qualitatively different from that of monolingual speakers. This was based on the observation in Carminati (2002) that monolingual speakers are more tolerant of violations of their preferred antecedents for overt pronouns in contexts with a single potential antecedent than in contexts with multiple potential antecedents. As such, Sorace (2011) suggested that bilingual speakers differ from monolingual speakers in that the former are more permissive of PAH violations for over pronouns, potentially due to cognitive load relating to bilingualism. However, such an account would only predict a relaxation of the non-SpecIP bias for overt pronouns under attrition. It could not capture the changes to null pronouns or PR/RC ambiguities. Moreover, even if the interpretation of overt pronouns was affected (depending on the reader's interpretation of the relevant interaction), this unambiguously did not surface as a relaxation of speakers' interpretive biases. In our experiment, the experimental group gave more, not fewer, topic-shift responses.

Our findings similarly do not fit nicely with Paradis's (1993, 2007) ATH. Recall that the ATH predictions regarding attrition rely on the notion of a competing form in the L2. It proposes that a linguistic element's activation (i.e., accessibility) for a speaker is related to its use (among other factors) such that the more recently and frequently a form is used, the more accessible it is whereas the less frequently or recently it is used, the less accessible it becomes. Under that account then, attrition is predicted to occur when a form in the L1 becomes less accessible than a competing form in the L2. However, as discussed above, our participants' L2 either lack an equivalent and therefore potentially competing form (in the case of null pronouns) or it does not make sense to argue that the form in the L2 is different in any relevant sense from the form in the L1 (in the case of the LA bias for RCs). It is also worth point out that in the thesis so far, we have tacitly made the simplifying assumption (which appears to also be present in Gürel 2004 et seq.) that the potential attriters L2-English representations are sufficiently similar to those of L1-English speakers (given their upper-intermediate to advance L2 proficiency) such that they do not analyse overt pronouns in English as structurally analogous to null pronouns in Italian (i.e., they are not an overt phonological realisation of *pro*). As a result, we have assumed that overt English pronouns only compete with overt pronouns in Italian. However, our argument does not hinges on this assumption. Rather it rests on the (admittedly untested) assumption that the potential attriters do not analyse overt subject pronouns in English as being specified for a feature like [±TS]. If that assumption is valid, it would not matter which of a speaker's L1 pronouns English pronouns are taken to compete with. Either way, one would expect a relaxation of interpretive biases, *contra* what was observed. However as all of the experiments in the present thesis only investigated the L1, we cannot evaluate any assumptions regarding their L2. Therefore, even though our results seem unexpected under the ATH, we cannot fully exclude the proposal and we leave this issue open to future research.

Comparison of the pronominal and PR/RC results from our sentence interpretation task also allows us to exclude another potential account for our PR/RC items that might otherwise seem salient. Given that we observed an overall shift toward more LA interpretations regardless of PR availability, one might interpret this as an increased susceptibility to intervention. That is to say, we might speculate that the increased LA bias is not due to participants relying more on their L1 biases as we have suggested, but rather due to them experiencing difficulty establishing a relationship across another intervening element of the same type, potentially due to the cognitive load associated with bilingualism (e.g., the need to maintain language control). However, if this were a general effect of attrition, this 'increased intervention' interpretation would make demonstrably false predictions regarding the pronominal items. Namely, it would predict an increased rate of non-SpecIP responses regardless of the pronominal form. This is because the non-SpecIP antecedents in our items always (linearly) intervened between the embedded pronoun and the antecedent in matrix SpecIP. Although there might have been an increase in the non-SpecIP interpretations for overt pronouns (depending on the reader's

interpretation of our interaction term), this is clearly not the case for null pronouns; there we observed a decrease, not an increase, in non-SpecIP interpretations. As such, if we tried to account for the PR/RC results in terms of intervention we would require an independent explanation for the pronominal data in addition to some sort of explanation for why the 'increased intervention' effect does not spill over into the pronominal items. As such, we would have to treat as accidental the generalisation that the experimental participants exhibited a stronger version of the bias expected from the control group in both of our phenomena. Therefore, an 'increased intervention' interpretation for our data should be dispreferred for reasons of parsimony.

In a similar vein, we can use the PR/RC results as additional evidence against some possible alternative interpretations of the pronominal data. Recall that based on Tsimpli et al.'s (2004) results with the same items, we predicted an increase in non-topic-shift responses for items with both null and overt pronouns. The results from our sentence interpretation task were partially consistent with this. We found evidence of an increased non-topic-shift bias for null pronouns and, at least under one interpretation of our interaction, no change in the items with overt pronouns. Nonetheless, previous studies including Tsimpli et al. (2004) and Chamorro et al. (2015a) have reported a relaxation of the topic-shift bias for overt pronouns in their experimental groups. As such, one might attempt to argue that pronouns in general begin to prefer more non-topic-shift interpretation under attrition with some caveat about why this was restricted to null pronouns in our version of the sentence interpretation task. However, such an 'increased topicality' interpretation cannot account for the PR/RC data. This is because neither of the relevant NPs in our items could be considered topical. As a result, although the 'increased topicality' interpretation cannot be fully excluded for the pronominal portion of our data, it cannot be extended to our other attrition effect. Therefore, similar to the 'increased intervention' interpretation, this would force us to say the apparent qualitative similarity between the attrition effects in our two phenomena is accidental. As such, the 'increased topicality' interpretation should be dispreferred for reasons of parsimony.

Thus having pointed out the various difficulties faced by the previously discussed processing-based explanations as well as some potential alternative interpretations of our data, it appears that we should take seriously the idea introduced chapter 2 that speakers undergoing attrition exhibit a general increased reliance on interpretive biases already found in their L1. Recall that we originally proposed this idea given that various previous studies reported what we argued should be interpreted as more extreme versions of biases already present in the L1. Moreover, this pattern reappeared in the interpretation of different forms (null pronouns: Tsimpli et al. 2004, Gürel 2004, Gürel & Yılmaz 2011; overt pronouns: Martín-Villena 2023; reflexives: Gürel 2004, Gürel & Yılmaz 2011; preverbal subjects: Tsimpli et al. 2004), suggesting that a phenomena-specific explanation would be inadequate. Applying this idea to our data has several advantages over the previously discussed alternatives. First, it is the only account so far discussed that would straightforwardly predict both the increase in LA interpretations for true RCs and the increase in non-topic-shift interpretations for null pronouns observed. As such, this proposal has greater empirical coverage than its alternatives. Second, a bias-couched account would directly encode the idea that the qualitatively similar way that attrition affects the two studied phenomena is not accidental. Thus, such an account is more parsimonious than positing a unique explanation for each of the affected phenomena such as if we assumed both the 'increased intervention' and 'increased topicality' interpretations discussed above. Third, such an account has the additional benefit of avoiding clashes where the phenomena-specific explanation for a given set of findings forces us to make incorrect predictions for another without further stipulation, as in the case of the 'increased intervention' interpretation and our pronominal data. In addition, it is worth noting that a bias-couched account does not require us to posit any additional theoretical machinery. This is because we must independently say that native speakers already possess these biases. Moreover, even if we assume that some of these biases such as the LA bias for RCs are universal as suggested by the PR Account, it is also clear that speakers' reliance on particular biases may be affected under certain circumstances. As a concrete example, this was directly evidenced by the adaptation effect observed in the online processing of overt pronouns in our online experiment. Thus

positing that speakers undergoing attrition rely more on their L1 biases seems like an attractive proposal. Nonetheless, to maintain such an account, we need to address some of our results that on the surface, do not appear to follow from it. Namely, under such an account the lack of evidence for an increased reliance on the non-SpecIP bias for overt pronouns or the PR-First Hypothesis in perceptual items would be unexpected. As such, to further explore whether our data is consistent with a bias-couched account let us now consider those two issues in turn.

As for how our overt pronoun results fit in with this 'increased biases' idea, this to some extent depends on the reader's preferred interpretation of our significant interaction in the offline experiment. However, in either case, the explanation is relatively straightforward and our results are fully compatible with the idea that our experimental participants exhibited an increased topic-shift bias. To see how this would work, let us suppose that the reader takes our relevant interaction to indicate that overt pronouns were also affected by attrition (albeit to a lesser extent than null pronouns). In that case, our results are straightforwardly compatible with the proposed interpretation. For both pronominal forms, the experimental participants exhibited a stronger, not a weaker, bias than the control group. Conversely, if we suppose that the reader interprets our significant interaction in the offline task to have been driven by a change in the items with null pronouns only, then it is worth highlighting that the responses from the control group were already relatively extreme for matrix-first items with overt pronouns. For those items, the control group selected SpecIP readings in only 12.88% of trials. This imposed a logical limit on any potential strengthening of the non-SpecIP bias given that percentage of SpecIP responses could not fall below 0. This issue is further compounded by the globally ambiguous nature of the items themselves. That is to say, it would be unreasonable to expect fully categorical judgments (i.e., 0% SpecIP) given that other responses were still possible. Thus, even if attrition were to have led to a strengthening of the non-SpecIP bias for overt pronouns in our experimental participants, any effect would have been subtle and difficult to detect with our items due to a floor/ceiling effect. (It is also worth pointing out that even if the reader prefers the other interpretation of our interaction term, this discussion of a floor effect may still go some way to explaining why the effect for overt was smaller than the effect for null pronouns.) Thus, regardless of which interpretation the reader prefers, our argument remains unchanged; we should not take the lack of an unambiguous effect in our items with overt pronouns to be incompatible with the 'increased biases' interpretation. This position is further supported by the fact that Martín-Villena (2023) reported the increased non-SpecIP interpretive bias for overt pronouns that would be predicted under this account in three distinct samples of now-bilingual native (Iberian) Spanish speakers.

As for why we did not observe an increased HA bias within the sentence-interpretation task's perceptual items, recall from chapter 3 that this could reasonably relate back to our discussion of why we observed an overall increased LA bias. Namely, all of our items that were compatible with a PR reading, were also trivially compatible with a RC reading. Moreover, we have good reason to believe that participants were accessing RC readings of those items given that (i) even control participants provided LA responses at a non-trivial rate and (ii) LA responses were incompatible with PR parses. As such, even if the participant exhibited an increased reliance on the PR-First Hypothesis, this would have conflicted with the increased locality effect, such that we were unable to detect the former in our sample. Given this potential experimental issue, we should not take this null result as a strong counterargument for a bias-couched account of attrition.

Having shown how our data could be captured under a bias-couched account, we should stress that we are not arguing that attrition must always lead to an increased reliance on biases already present in the L1. That position would be too strong and is demonstrably false. Focusing on overt pronouns, there is non-trivial previous evidence that the interpretive and processing biases for such forms in null-subject languages may weaken under immersion in a non-null-subject L1 as discussed in chapter 2 (e.g., Italian: Tsimpli et al. 2004; Turkish: Gürel 2004; Spanish: Chamorro et al. 2015a; Greek: Kaltsa et al. 2015). Moreover, at least in Turkish, this effect has been shown to be replicable with a different set of participants (Gürel & Yılmaz 2011). Thus even though we were not able to replicate the overt pronoun result from Tsimpli et al.

(2004) in our sentence-interpretation task, the fact that various other researchers have found related effects suggests that Tsimpli et al.'s (2004) effect is unlikely to have been spurious. As such, it seems we are forced to say that relaxation of biases does occur as would be predicted by accounts like Sorace's (2011, 2016) overt-as-processing-default idea or Paradis's (1993, 2007) ATH.

The effect for overt pronouns clearly does not follow from the kind of bias-couched account discussed above and raises the question of how we could modify such an account in light of this finding. A potential response might be to suggest that we need a mixed account, i.e., a bias-couched account plus the reader's preferred account for the overextension of overt pronouns. This seems possible at least in principle because a biases-couched account is not mutually exclusive with the above-mentioned processing-based accounts. That is to say, it may well be the case that speakers undergoing attrition tend to rely more on their L1 biases in general but this tendency is overridden in specific cases by a tendency to overextend certain forms due to a default processing form or a competing element in the L2. However, if we reject the idea that attrition should receive a unified explanation, this clearly comes at the cost of parsimony. Potentially more problematically, this would also force us to make conflicting predictions. To see why this is the case, assume for a moment that we adopt the idea that overt pronouns function as a processing default. While a bias-couched account would predict a strengthening of the non-SpecIP bias for overt pronouns, an overt-as-processing-default account would at the same time predict a relaxation. Although the reader's immediate reaction to such conflicting predictions may be understandably less than positive, this state of affairs may be necessary based on the currently available data. Namely, there seems to be inconsistency in the way that attrition surfaces for overt pronouns which is exactly where we would expect the two tendencies discussed above to conflict. On the one hand, there is compelling evidence that the non-specIP bias for overt pronouns may weaken under attrition. On the other hand, we cannot disregard Martín-Villena's (2023) results which indicated a strengthening. Recall that they observed a strengthening in three distinct groups of participants which were noticeably larger (e.g., immersed experimental group N = 94) than those in other studies on attrition (e.g., Tsimpli et al. 2004: Italian attriters N = 20; Chamorro et al. 2015a: experimental group N =24).

Assuming that the reader is willing to entertain the possibility of a mixed account, we should also highlight that such a model would face a non-trivial challenge. Namely, without some further restriction on when we should expect a weakening or a strengthening of biases when the two tendencies conflict, a mixed account of attrition would be unfalsifiable in certain contexts. For example, any possible outcome would be predicted for overt pronouns. Therefore, a fully unrestrained mixed account of attrition is untenable. If we are to take this idea seriously, we must somehow restrict the model such that we can make falsifiable predictions about when we should see a strengthening or a weakening.

Although the present thesis has largely focused on the global comparison of pronominal reference and PR/RC parsing ambiguities, a potential way to resolve the observed inconsistencies in the pronominal data across studies might be to recognise that there also is to a certain extent variation within the two phenomena investigated. For example, within the PR/RC items, even if we assume that the PRFH and Late Closure (LC) are universal parsing principles driven by computational efficiency at the level of the parser, arrival at a PR parse for a one of our PR-compatible string still requires the successful integration of the selectional properties of the matrix predicate. In the case of a true RC attachment ambiguity however, the integration of the matrix predicate is not required. Turning to the pronominal items, we might follow the Form-Specific Multiple Constraint approach proposed for Finnish by Kaiser & Trueswell (2008). That approach suggests that pronominal resolution is affected by multiple factors and that these multiple factors may affect specific pronominal forms differently. Transferring this idea to Italian (and Spanish), recall that in section 2.2 we discussed Carminati's (2002) experiment with a single textually given antecedent. Results form that experiment suggested that while the SpecIP bias for null pronouns is unaffected by context, the non-SpecIP bias for overt pronouns is more flexible and sensitive to other contextual factors. This raises the possibility that the conflicting results for overt pronouns between studies may relate to the different set of items/tasks employed. Some

initial credence for this idea can be taken from Martín-Villena's (2023) who reported very similar increased non-SpecIP bias for overt pronouns in 3 distinct samples of Spanish speakers using the same stimuli and task. Nonetheless, given the currently available results we cannot tease out what specific factors might play role in the conflicting results for overt pronoun, and so this idea remains speculation. The only potential restriction on a bias based account that we can easily judge based on the current data is the idea that the more specific tendency (i.e., overextention) always trumps the more general tendency (i.e., to rely more on the L1 biases). Given the results from Martín-Villena (2023), this is clearly wrong. Beyond this, we leave the question of how to restrict such a model open to future research. However, as a final point on this empirical complication, it is worth pointing out by itself this should not be taken as an argument for dispreferring an account that makes reference to biases over the alternative processing-based accounts. Recall that we have discussed at length that models such as Sorace's (2011, 2016) overt-as-processing-default or the Paradis's (1993, 2007) ATH face empirical problems of their own in that they cannot account for the fact that biases for a number of other phenomena appear to strengthen.

As a final point before integrating our online results, it is also worth noting that a bias-couched (or mixed) account does not require us to abandon the 'selectivity of attrition.' An account along these lines would still be able to capture the fact that the licencing of null subjects (and related properties such as the absence of *that*-trace effects) is unaffected by attrition whereas the interpretation of null and overt pronominal subjects is. Moreover, we can also still capture the idea expressed by more recent instantiations of the IH that structures at more internal interfaces are generally less affected than structures at more external interfaces, under the assumption that the interpretation of structures at more external interfaces tends to be less categorical. Specifically, take Chamorro et al.'s (2015b) study on DOM, which she argues to be a syntax-semantics interface structure as the presence of the personal proposition ('a') marking a direct object (DO) is conditioned by factors including the animacy of the DO. They observed that DOM is stable[79] under attrition despite their participants exhibiting clear attrition effects in their processing of pronominal forms during the same experimental session (Chamorro et al. 2015a). The original authors take that asymmetry as evidence that structures at the interface with semantics are more protected from attrition than structures at interface with discourse. However, the asymmetry in their results is also fully consistent with the proposal in the present thesis even though it does not make reference to interfaces. That is because the interpretative preferences for null and over pronouns are violable, but Chamorro et al. (2015b) specifically designed their stimuli such that the acceptability of DOM was functionally binary; either the presence/absence of DOM was fully acceptable, or it was fully unacceptable. As such, there was no room for any increased reliance on the L1 biases to manifest.

Turning now to our experiments investigating online processing, the results from the overt pronominal items are also convergent with the 'increased biases' interpretation for the offline data. Recall that toward the beginning of the self-paced reading experiment, the participants living abroad exhibited a more pronounced effect of antecedent (i.e., non-SpecIP items were read more quickly than SpecIP ones) relative to the participants still living in Italy, with the caveat that this difference diminished as the experiment progressed. As with the increased SpecIP bias for null pronouns and the increased LA bias for the RCs in the sentence interpretation task, this increased non-SpecIP bias for overt pronouns in the experimental group (toward the beginning of the experiment) cannot be attributed either to transfer or to overt pronouns functioning as a processing default. Starting with the possibility of transfer, participants were L2 speakers of English. Given the non-null-subject status of that language, we follow Tsimpli et al. (2004) in assuming that overt subjects in the speakers' L2 are not specified for a feature like [±TS]. As such, there is nothing to transfer that would somehow reinforce the L1 bias. If anything, were the speakers to have transferred properties from their L2 we would have expected a relaxation of the non-SpecIP bias along the lines of Tsimpli et al.'s (2004) proposal.

[79] Although, see also Montrul et al. (2015) for conflicting results on DOM from another variety of Spanish.

Moreover, our result again cannot be accounted for by appealing to overt pronouns as a default form during processing or the ATH. Both accounts would predict that the penalty for violating the antecedent preference for overt pronouns would be reduced in the experimental group, resulting in a weakened non-SpecIP bias. By extension, this would reverse the direction of our three-way interaction, contrary to what we observed. As such, we take our overt pronoun results from chapter 5 to suggest that attrition can not only lead to an increased reliance on interpretive biases at the level of the final interpretation but also that this increased reliance on the L1 biases extends to online processing.

This result and our interpretation of it clearly contrast with those in Chamorro et al. (2015a). Recall that they reported significant *pronoun* by *antecedent* interaction in the eye-tracking measures from their control group of native (Iberian) Spanish speakers which was driven by slower reading times for overt pronouns when they were forced to co-refer with the potential antecedent in SpecIP. However, no such significant interaction was observed in the results from the (non-re-exposed) experimental group, resulting in a significant three-way interaction when the two groups were compared against each other. They interpreted this to indicate that native (Iberian) Spanish speakers undergoing attrition are less sensitive to their L1 biases in online processing. Nonetheless, as discussed above, our result is not without precedent and there seems to be some inherent inconsistency in the way that overt pronouns in null-subject languages are affected by attrition that is not yet understood. As such, we should not take the conflict with the results Chamorro et al. (2015a) as strong evidence against our conclusion based on our own data that attrition may sometimes lead to a strengthening of the non-SpecIP bias even in online processing.

Comparing the pronominal results from the self-paced reading experiments with those from the sentence interpretation experiment raises another question about our pronominal results in general. Namely, why should one form be affected in offline judgments while the other is affected in online processing? We suggest that this asymmetry is likely only apparent and could reasonably relate to the tasks themselves. To see why this is the case, let us consider each pronoun in turn.

Starting with overt pronouns, we have already discussed at length that regardless of the reader's preferred interpretation of the relevant interaction, our offline results are consistent with an 'increased biases' interpretation.

Turning to null pronouns, the lack of an attrition effect in our online experiment might lead a sceptical reader to question whether we should interpret the results from our offline task. However, as pointed out above, the fact that our experimental group exhibited an increased SpecIP bias for null pronouns in the sentence interpretation task closely replicates what was reported in Tsimpli et al. (2004) suggests that our offline effect is unlikely to have been spurious. Instead, we take it to be more reasonable that the lack of an attrition effect in the online experiment relates to an experimental artefact in that experiment instead. In fact, we independently argued that our items may not have been ideal specifically for detecting the SpecIP bias for null pronouns even before encountering the data from the experimental group. This is because we were unable to detect a SpecIP bias in the processing of the null pronouns in either of our online experiments that only considered native Italian speakers still living in Italy. This is notable given that previous authors have repeatedly found a SpecIP bias with different sets of items and participants (Carminati 2002, 2005, Vogelzang et al. 2019) and some of those results have been demonstrated to be replicable (Filiaci et al. 2014). Given the absence of this baseline effect when we have good reason to expect it, we take the lack of a significant interaction with *group* in Experiment 3 to be uninterpretable. As such, we leave it to future research with different items to investigate whether the increased SpecIP bias for null pronouns might extend to online processing.

As an additional point regarding the pronominal items in our online experiments, now that we have seen the results from the PR/RC items in the same experimental sessions, we may exclude a salient potential explanation for the lack of the expected baseline effect. Namely, we can exclude the possibility that the asymmetry in our online pronominal results is related to the fact that the disambiguating gender marking surfaced differently

in the items with null and overt pronouns. In the case of a null pronoun, the disambiguating gender marking only surfaced on the secondary predicate within the embedded CP. In the case of an overt pronoun, however, the gender marking was also manifest on the pronoun itself. As such, if we only considered the pronominal results, we might suggest that the gender marking on the embedded secondary predicate was insufficiently salient during their initial online processing such that participants were only able to integrate this information at a later stage. Such an account would capture why we did not observe an online effect specifically for null pronouns as well as participants' high accuracy to comprehension questions overall. However, this would make patently false predictions for the PR/RC items. Namely, we would have expected no online effects in those items, given that they were all disambiguated exclusively via gender marking on the secondary predicate. Thus again on grounds of parsimony, we should reject this potential explanation for the lack of an online effect within the null pronominal items.

Although integrating the PR/RC items from the online experiment can help us rule out a potential interpretation of our pronominal items, it also introduces another potential complication for our discussion. Namely, if attrition affects syntax-discourse interface structures and attachment ambiguities similarly as we argued above, then it seems unexpected that we would find an attrition effect in the online processing of our pronominal items but not the PR/RC items. This assumption seems particularly sensible given the high level of overlap in the participant samples for those two experiments. Thus we might interpret this asymmetry to indicate that attrition is selective and only affects the online processing of PAH biases but not the PR-First Hypothesis or the LA bias for true RCs. However, we should be extremely cautious of this interpretation for several reasons. First, it rests on a null result. Second, it would require us to somehow explain why the LA bias for RCs may be affected in offline interpretations, but not in online processing. Third, by extension, it would likely force us to abandon a unified account for the increased biases observed in the interpretation of pronominal and PR/RC ambiguities despite the discussion above. Fourth, taking seriously the idea that attrition does not affect the online procession of PR/RC ambiguities would require us to discount the significant and consistent results in Dussias (2003, 2004) and Dussias & Sagarra (2007). Given these issues, we should not take the lack of significant interaction with *group* in our online PR/RC study as strong evidence that attrition does not affect these structures. As such, we leave the question of how attrition affects these parsing of PR/RC ambiguities (RQ_8) to future research.

Thus, taking stock of the discussion so far, we have been primarily concerned with establishing (i) how attrition affects the interpretation and processing of PR/RC ambiguities, (ii) whether those effects were qualitatively similar to the effects observed for pronominal ambiguities, and (iii) how we should best interpret these effects. Our response has been based largely on our offline results and we have argued that attrition affects these two phenomena in a qualitatively similar way. Moreover, we argued that this should be interpreted as a bilingualism effect (i.e., it is not due to transfer), and the observed patterns cannot be accounted for by previous processing-based accounts. As such, we have proposed that our data may be captured by positing that speakers undergoing attrition may begin to rely more heavily on their L1 biases and have discussed some tension that arises when we consider such a bias-couched account and the overextension of overt pronouns reported by other authors. So far, however, we have not discussed why speakers might rely more on their L1 biases as a result of becoming bilingual. This is a necessary but more difficult question and one which the studies in the present thesis were not designed to explore. Nonetheless, as a starting point, we might still consider whether/how some of the more cognitive explanations for the IH might be applied to our data and a bias-couched account instead.

In terms of real-time processing, Sorace (2011) points out that the felicitous use of pronominal forms requires the rapid integration of grammatical and discourse information. For example, speakers must update their mental representations of the discourse as the context changes, assess what information is shared with their interlocutor, and exclude inappropriate pronominal mappings to potential antecedents. As a result, they suggested that the particularly demanding nature of the phenomena in question combined with the plausible

assumption that bilingual speakers might be susceptible to cognitive load leads bilingual speakers to resort to overt pronouns as a processing default more so than monolingual speakers. Although we have discussed at length that overt pronouns serving as a processing default cannot capture the attrition effects found in the present thesis, we might wonder if we can extend the idea that the demands of the phenomena themselves might have contributed to the attrition effect. However, such an approach seems infeasible for our RC results given that the resolution of attachment ambiguities in our experiments did not require participants to consider the discourse context. That is to say, we have no reason to believe the interpretation of such forms is particularly demanding akin to the interpretation of pronominal forms. As such, if try to maintain this idea for our pronominal data, we would be forced to treat the qualitative similarity between how attrition affected our pronominal and PR/RC items as accidental.

Instead, we might want to explore the related idea of overexplicitness. For this, Sorace (2016) builds Hendriks, Koster & Hoeks's (2014) model of reference tracking. In that model, the choice of referring expression during production depends on the speaker (i) first selecting the most reduced form (e.g., null pronouns) and (ii) then selecting an alternative form that can most easily be understood by their listener if necessary. As such this model predicts that the production of less reduced forms (e.g., overt pronouns) is more costly as it requires the speakers to consider the listener's potential interpretation. Assuming that model, then Sorace (2016) suggests we might re-interpret the overextension of overt pronouns (in production) not as evidence of cognitive load, but rather as bilingual speakers avoiding more reduced forms, potentially due to an increased ability to consider the perspective of the interlocutor. Nonetheless, if we again consider comprehension, Sorace (2016) suggests that speakers initially interpret referential forms as mapping onto the most prominent potential antecedent. In the case of null pronouns, that is unproblematic. In the case of overt pronouns, however, overcoming that initial prominence bias requires the speaker to consider the mental state of their interlocutor, resulting in the previously observed interpretive overextension. Setting aside some of the issues this idea faces with regard to the pronominal data (e.g., how is it that speakers exhibit 'enhanced perspective-taking abilities' during production but at the same time face difficulties considering a speaker's intentions during comprehension?), this idea is extremely similar to our 'increased topicality' interpretation we considered for the sentence-interpretation task results and runs into the same problem. Namely, such an account can again not be generalised to capture the increased LA bias we observed for true RCs, and as such should be dispreferred.

As an alternative, we might follow Sorace (2011, 2016) in considering the potential role of cognitive load due to bilingualism itself. Namely, as our experimental participants were all proficient bilinguals living in their L2 community at the time of testing, it is reasonable to assume that their L2s were highly activated. Therefore if late L2ers rely on inhibition to maintain language control (e.g., Costa & Santesteban 2004) then it might be the case that language control and the processing of referential expressions draw from the same set of cognitive resources. Assuming that is the case for a moment, competition for these resources might then induce cognitive load that sometimes exceeds a speaker's resources. In Sorace's terms that might lead to the discoordination of cognitively demanding pronominal processing. That in turn would result in bilinguals' overextension of overt pronouns either due to overt pronouns functioning as a processing default or because participants are less able to consider their interlocutors' mental states. Recasting this idea in terms of biases, if we were to speculate that overcoming one's initial interpretation/parse required inhibition, then we might suggest that our experimental participants were less able to consider alternative interpretations/parses. This combined with the fact that our control participants were far from categorical in their biases, may have resulted in the apparent strengthening of the L1 biases. Such an idea is tempting for our sentence-interpretation data as it also would straightforwardly apply to both our pronominal and PR/RC items. However, a cognitive explanation along these lines would make at least two predictions for that task, one of which immediately appears to be wrong. First, as pointed out by Sorace (2016) the size of the observable bilingualism effect would vary with the difficulty of maintaining language control. That is to say, if we re-conducted our experiment

with L1-English L2-Italian speakers, we might expect similar but more pronounced bilingualism effects under the assumption that for most speakers the inhibition of the L1 is likely more demanding than the inhibition of the L2 as suggested by asymmetric switching costs (Meuter & Allport 1999, although see also Gade, Declerck, Philipp, Rey-Mermet & Koch 2021 and Goldrick & Gollan 2023 for recent meta-analyses of that literature). Second, the size of the observable effect would vary with individual differences in executive function. Those speakers who are more able to inhibit irrelevant information might be expected to exhibit fewer bilingualism effects. While the experiments presented in the present thesis did not consider inhibition and our narrow focus on attrition meant that we did not test L2 speakers of Italian, there is nonetheless some evidence against the first of these predictions. Both Sorace & Filiaci (2006) and Belletti et al. (2007) conducted Tsimpli et al.'s (2004) sentence-interpretation task with L1 English participants who had achieved 'near native' proficiency in their L2 Italian as well as native Italian controls. Although the values from their control groups were comparable to those in Tsimpli et al. (2004) and the present thesis, in neither L2 study do the authors find a difference between their experimental and control groups with regard to the SpecIP bias for null pronouns. This is clearly inconsistent with the prediction that L2 speakers would exhibit a larger bilingualism effect as a result of greater cognitive load to maintain language control.

As an alternative to those more cognitive explanations, we might instead suggest that attriters' increased biases relate to their metalinguistic awareness about the differences between their L1 and their L2. That is to say, it is possible that upon realising that their L2 lacks null subject pronouns of the type found in the L1, they might also realise the interpretive differences between overt pronouns in their two languages. As a result, in an experimental context, they may rely more on that explicit knowledge than their linguistic intuitions, resulting in more extreme biases. However, there are at least two salient reasons not to pursue this idea. First, it seems reasonable that such an account would predict attrition to mostly affect the interpretation of globally ambiguous elements and would not affect their online processing. Yet, we detected an increased reliance on the non-SpecIP bias in the online processing of overt pronouns. Second, tying attrition effects to metalinguistic awareness again treats the similarities between the attrition effects in our pronominal and PR/RC items as accidental. After all, speakers may realise that their L2 lacks null pronouns but without realising it also lacks PRs or *vice versa*. Thus, although we have argued that there are good reasons to interpret our data in terms of increased reliance on one's L1 biases, it is not clear how we could explain why this is the case. As such, we leave this issue open to future research.

A final point that should be discussed with regard to potential accounts for attrition is the fact that the present thesis opted for a between- rather than within-group experimental design. Ultimately this decision was motivated by our research questions. We were interested in comparing how attrition affected different phenomena *within* the same speakers (or very similar sets of speakers in the case of chapters 5 and 6) rather than investigating factors that predict attrition *between* participants. As such, our group-based design was appropriate. However, this came with both advantages and limitations. On the one hand, this allowed for more direct comparability to previous samples of attriters (e.g., Tsimpli et al. 2004, Chamorro et al. 2015a), while also reducing the need for arbitrary choices in our modelling that were not directly dictated by research questions / hypotheses (e.g., instead of *group*, should the model included *length of residency*, *L2 proficiency*, or some index of dominance index among others, and if the model contained a continuous predictor such as *length of residency* should the effect be modelled linearly?). In addition this avoided inflating the risk of type 1 error due to the inclusion of multiple potential predictors (e.g., if model included both *length of residency* and *L2 proficiency*, would a significant interaction with either count as finding an effect of attrition?). On the other hand, our relatively strict inclusion criteria for the experimental group and use of a control (but not necessarily monolingual) baseline means that our results may not be representative of (L1-Italian, L2-English) bilinguals more generally. Additionally, our strict inclusion criteria limited our ability to use individual differences within the experimental participants to explore which aspects of their language background (e.g., L2 exposure, L2 proficiency, exposure to local varieties as children) or cognitive profiles (e.g., executive functions, or reliance

on inhibition to maintain language control) are most relevant to attrition as was done in Martín-Villena (2023) for example. In turn, this meant that we could not use individual differences to further evaluate potential accounts of attrition.

Moving beyond attrition, in order to explore our research questions, the present thesis also presented a number of results that are relevant for researchers interested in the PR Account. Starting with our sentence-interpretation task (chapter 3), we demonstrated that the results in Grillo & Costa (2014) are straightforwardly replicable. Despite Italian being a so-called HA language, participants exhibited a LA bias for true RCs and an HA bias only appeared where PRs were locally available as would be predicted by the PR account. Moreover, that experiment provided evidence that the bias for PRs over RCs is stable under attrition despite the global shift toward more LA interpretation. That is a welcome result if the PR-First Hypothesis is taken to be a universal bias. Moreover, results from our online tasks were able to validate a number of predictions derived from the PR Account. Namely, when PRs are locally unavailable, native speakers of Italian again repeatedly exhibited a LA bias. However, when PRs were locally available, we observed HA bias at the post-disambiguating window as has previously been reported for native speakers of Italian (De Santo & Lee 2022). In addition, our self-paced reading task with speakers still living in Italy (chapter 6) also extended this to the disambiguation point itself consistent with previous results from Spanish (Aguilar et al. 2021). Moreover, we have presented a novel type of evidence in support of the PR Account. Namely, in all three of our experimental samples, we were able to detect an influence of the PR-First Hypothesis in the interpretation of fully disambiguated items. In all of the relevant analyses, participants were significantly less accurate when responding to LA-disambiguated perceptual items than HA-disambiguated ones.

Although it was not our primary focus, the effect in the comprehension questions for our PR/RC items is also likely to be of unexpected interest to researchers working on persistence effects for several reasons. First, although we were able to detect the persistence of intermediate PR parses, in the same participants, we found no evidence that responses to HA-disambiguated non-perceptual items were more accurate than to LA-disambiguated ones. That is to say, there was no persistence effect for RC. This was despite clear and repeated effects of locality in many of our online measures for the same items and resulted in a significant interaction when the responses to perceptual and non-perceptual items were modelled together. This raises the question of how we should account for this asymmetry because regardless of whether persistence effects are due to difficulties in reanalysing the string (Christianson et al. 2001) or in erasing an initial interpretation (Slattery et al. 2013), the asymmetry is unexpected. Moreover, given that this asymmetry appears to be replicable in different samples of native Italian speakers, it appears it should not be considered spurious. This suggests that it is not enough for participants to be garden-pathed in order to observe persistence effects. Rather, we suggested in chapter 6 that there might be something particularly difficult in the reanalysis of the matrix DO. As such, future work on persistence effects should investigate whether this asymmetry generalises to other contexts and, if so, explore how we might account for this.

As a second point regarding persistence effects, in none of our datasets, did the best fitting model include an interaction term with *response target*. This is worth highlighting as Tabor, Galantucci & Richardson (2004) suggested that the persistence effect originally observed in items like (87) by Christianson et al. (2001) may have been due to their comprehension questions reactivating the intermediate parse. If such a counterargument were true, this would lead to a prediction for our own data given that we counterbalanced our questions by the response target. Namely, it would predict that where the intermediate parse for some item conflicted with the globally correct parse, responses would have been significantly more accurate when the target response was *sì* ('yes'). As such, we would have expected a significant three-way interaction with *response target*. However, in none of our analyses, even the one over the pooled data from Experiment 1 and Experiment 3 with a total of 98 participants, did our model selection process via warranted inclusion indicate that there was sufficient evidence for an interaction with *response target*. As such, it seems highly improbable that our result could have been driven by reactivation (i.e., 'no' questions). This interpretation is convergent with work that has

detected persistent effects using designs that explicitly avoid repetition (e.g., van Gompel, Pickering, Pearson & Jacob 2006, Malyutina & den Ouden 2016) and suggests that the effect should not be interpreted as an artefact of the experimental design.

(87) While the man hunted the deer ran into the woods.

Q. Did the man hunt the deer?

Third, our results from the experimental group in Experiment 3 may also be relevant for work on L2 reanalysis. This is because some previous authors have used persistence effects to investigate how successful reanalysis is in the L2 compared to in the L1 (e.g., Fujita & Cunnings 2021, Jacob & Felser 2016, Pozzan & Trueswell 2015). Exemplifying this Jacob & Felser (2016) presented sentences like that in (88) to a group of native English speakers as well as a group of native German speakers who had achieved 'upper intermediate' proficiency in their L2-English.

(88) While the gentleman was eating the burgers were still being reheated in the microwave.

Q. Was the gentleman eating the burgers? Yes or no?

In such sentences, the DP immediately after the lexical verb of the embedded CP ('the burgers') may be initially parsed as the DO of the embedded lexical verb ('eating'). However, upon encountering the tensed element in the matrix CP ('were'), the preceding DP must be reanalysed as the matrix subject. As such, when presented with the comprehension question in (88Q), the expected answer would be 'no' if participants fully reanalysed the sentence. However, if the intermediate DO parse persisted, participants might answer 'yes' either because they were not able to fully reanalyse the item or due to interference from the intermediate parse. Consistent with Christianson et al.'s (2001) original work, the native English speakers were observed to respond correctly (i.e., no) in only 72% of critical trials despite 89% accuracy in unrelated filler items. More germane to the present discussion, the effect was significantly larger in the L2 group who responded correctly only in 54% of trials, despite their response accuracy for filler items being comparable to the L1 group (88%). This has been taken as evidence that L2 speakers may be less able to revise their initial parse or remove it from their working memory. The fact that we were unable to detect a similar difference between our control and experimental groups may suggest that the L2ers increased susceptibility to an interfering initial parse cannot be explained in purely cognitive terms. This is because any cognitive demands imposed by bilingualism (e.g., due to the need to maintain language control) would be shared by L2 speakers and potential attriters. Of course, we should be cautious of the reasoning for two reasons. First, work on persistence effects in L2 parsing has focused on intermediate or intermediate-to-advanced speakers (e.g., Fujita & Cunnings 2021). As such, it might be that such effects become attenuated at the highest levels of proficiency (i.e., in 'near-natives'). Second, our reasoning relies on a null result. As such, this idea should not be given great importance by itself. However, our items and data are publicly available. Therefore, future work could straightforwardly reconduct our study with a population of L1-English speakers who have become highly proficient in their L2-Italian. If there truly is something qualitatively different about L1 and L2 reanalysis such that L2ers are more susceptible to persistence effects, then we would predict that the L2 group differs from both the control group and the experimental group in the present thesis, despite both experimental groups consisting of proficient bilinguals.

7.4 Conclusion

In conclusion, the present thesis has argued that although attrition may lead to the relaxation of certain biases, it may also lead to an increased reliance on the biases in one's L1 more generally. This was empirically demonstrated in both globally- and locally-ambiguous contexts. Moreover, we have presented evidence that

this pattern is not restricted to structures at the interface with discourse but also affects RC attachment ambiguities in a qualitatively similar way, suggesting that any structure which is interpreted non-categorically but subject to an interpretive bias may be affected regardless of interface status. As the attrition effect was not attributable to transfer from the L2 in either of the structures considered, we have argued that our results should be interpreted in terms of a bilingualism effect and as such should receive a unified account. This lead us to consider various alternatives way of achieving this, from which we argued that the data is best captured by an account that is (at least partially) couched in terms of the biases themselves.

A Data accessibility

Data and R scripts for the results reported in the thesis are available at the following links.

Chapter 3 Attrition and interpretive biases in globally ambiguous items
Link: https://osf.io/7kdmr/?view_only=79ea5d00af6b4d44a31dbe47083ab4fc

Chapter 5 Attrition and the Position of Antecedent Hypothesis in locally ambiguous items
Link: https://osf.io/fdq2m/?view_only=beef5c0f78a946789c0ab18dd8737769

Chapter 6 Attrition and the PR-First Hypothesis in locally ambiguous items
Link: https://osf.io/gr3yh/?view_only=c7174e3c5e4c4720b92ee8ae9de27b37

B Globally ambiguous stimuli

Num.	Predicate	Sentence
1	perceptual	Gianni ha visto il figlio del medico che correva la maratona.
	non-perceptual	Gianni vive con il figlio del medico che correva la maratona.
		'Gianni (saw / lives with) the son of the doctor that was running the marathon.'
2	perceptual	Maria ha sentito la nonna della ragazza che gridava.
	non-perceptual	Maria lavora con la nonna della ragazza che gridava.
		'Maria (heard / works with) the grandmother of the girl that was screaming.'
3	perceptual	Pietro ha sentito il maestro del ragazzo che cantava.
	non-perceptual	Pietro si allena con il maestro del ragazzo che cantava.
		'Pietro (heard / works out with) the teacher of the boy that was singing.'
4	perceptual	Lo scrittore guardava la zia della ragazza che saltava.
	non-perceptual	Lo scrittore ha sposato la zia della ragazza che saltava.
		'The writer (watched / married) the aunt of the girl that was jumping.'
5	perceptual	Silvia ascoltava la figlia del poliziotto che parlava.
	non-perceptual	Silvia lavora per la figlia del poliziotto che parlava.
		'Silvia (listened to / works for) the daughter of the policeman that was talking.'
6	perceptual	Paola osservava l'amico del politico che cucinava.
	non-perceptual	Paola è fidanzata con l'amico del politico che cucinava.
		'Paola (observed / is in a relationship with) the friend of the politician that was cooking.'
7	perceptual	Mario ha sorpreso l'assistente dell'attrice che rubava.
	non-perceptual	Mario è affezionato all'assistente dell'attrice che rubava.
		'Mario (surprised / is fond of) the assistant of the actress that was stealing.'
8	perceptual	L'avvocato ha beccato l'autista del vicino che fumava.
	non-perceptual	L'avvocato si esercita con l'autista del vicino che fumava.
		'The lawyer (caught / works out with) the driver of the neighbout that was smoking.'
9	perceptual	Lucia osservava il vicino del segretario che si allenava.
	non-perceptual	Lucia è innamorata del vicino del segretario che si allenava.
		'Lucia (observed / is in love with) the neighbour of the secretary that was working out.'
10	perceptual	Giorgio guardava il nipote dell'infermiera che mangiava.
	non-perceptual	Giorgio è imparentato col nipote dell'infermiera che mangiava.
		'Giorgio (watched / is related to) the nephew of the nurse that was eating.'
11	perceptual	Carlo ha fotografato il collega dell'impiegato che rubava.
	non-perceptual	Carlo odia il collega dell'impiegato che rubava.
		'Carlo (photographed / hates) the colleague of the worker that was stealing.'
12	perceptual	Sara ha visto l'amico del giudice che guidava.
	non-perceptual	Sara convive con l'amico del giudice che guidava.
		'Sara (saw / lives with) the friend of the judge that was driving.'

continued on the next page

Num.	Predicate	Sentence
13	perceptual	Francesco immaginava l'amica dell'estetista che lavorava.
	non-perceptual	Francesco cena con l'amica dell'estetista che lavorava.
		'Francesco (immagined / dines with) the friend of the hairdresser that was working.'
14	perceptual	Rachele ha sognato l'amico del cugino che beveva.
	non-perceptual	Rachele è sposata con l'amico del cugino che beveva.
		'Rachele (dreampt / is married to) the friend of the cousin that was drinking.'
15	perceptual	Ennio ha ritratto il fratello della donna che fumava.
	non-perceptual	Ennio lavora per il fratello della donna che fumava.
		'Ennio (dipicted / works for) the brother of the woman that was smoking.'
16	perceptual	Filippo ha filmato l'agente del giocatore che russava
	non-perceptual	Filippo frequenta l'agente del giocatore che russava
		'Filippo (filmed / dates) the agent of the player that was snoring.'
17	perceptual	Maria ha registrato il cugino dell'avvocato che parlava.
	non-perceptual	Maria lavora per il cugino dell'avvocato che parlava.
		'Maria (recorded / works for) the cusin of the lawyer that was talking.'
18	perceptual	Roberto ha guardato l'amico del pizzaiolo che ballava.
	non-perceptual	Roberto ama l'amico del pizzaiolo che ballava.
		'Roberto (watched / loves) the friend of the policeman that was dancing.'
19	perceptual	Simona ha fotografato il vicino dell'infermiera che studiava.
	non-perceptual	Simona collabora col vicino dell'infermiera che studiava.
		'Simona (photographed / collaborates with) the neighbour of the nurse that was studying.'
20	perceptual	Michele guardava il fratello del manager che scalava.
	non-perceptual	Michele studia col fratello del manager che scalava.
		'Michele (watched / studies with) the brother of the manager that was climbing.'
21	perceptual	Antonio ha filmato la sorella dell'amica che scriveva.
	non-perceptual	Antonio ha sposato la sorella dell'amica che scriveva.
		'Antonio (filmed / married) the sister of the friend that was writing.'
22	perceptual	Mario immaginava l'amica della collega che ballava.
	non-perceptual	Mario lavora con l'amica della collega che ballava.
		'Mario (imagined / works with) the friend of the colleague that was dancing.'
23	perceptual	Massimo ha visto l'insegnante dell'amica che guidava.
	non-perceptual	Massimo esce con l'insegnante dell'amica che guidava.
		'Massimo (saw / goes out with) the teacher of the friend that was driving.'
24	perceptual	Anna ascoltava il figlio del vicino che cantava.
	non-perceptual	Anna studia col figlio del vicino che cantava.
		'Anna (listened to / studies with) the son of the neigbour that was singing.'

Table B.1: PR/RC items borrowed from Grillo & Costa's (2014) second experiment with added translations.

Num.	Direction	Pronoun	Sentence
1	Mat-first	null	L'anziana signora saluta la ragazza, quando ø attraversa la strada.
		overt	L'anziana signora saluta la ragazza, quando lei attraversa la strada.
			'The elderly woman greets the girl, when she crosses the street.'
2	Mat-first	null	La segretaria aitua l'infermiera, mentre ø scrive una lettera.
		overt	La segretaria aitua l'infermiera, mentre lei scrive una lettera.
			'The secretary helps the nurse, while she writes a letter.'
3	Mat-first	null	Il nonno parla al nipote, mentre ø legge un libro.
		overt	Il nonno parla al nipote, mentre lui legge un libro.
			The grandfather speaks to the grandson, while he reads a book.'
4	Mat-first	null	La nonna mostra una foto alla nipote, mentre ø sta mangiando.
		overt	La nonna mostra una foto alla nipote, mentre lei sta mangiando.
			'The grandmother shows a photo to the grandaughter, while she is eating.'
5	Mat-first	null	La ragazza bionda dà i documenti all'impiegata, appena ø entra null'ufficio.
		overt	La ragazza bionda dà i documenti all'impiegata, appena lei entra null'ufficio.
			'The blond girl gives the documents to the employee, as soon as she walks into the offcie.'
6	Mat-first	null	La mamma dà un bacio alla figlia, mentre ø si mette il cappotto.
		overt	La mamma dà un bacio alla figlia, mentre lei si mette il cappotto.
			'The mother gives her daughter a kiss, while she puts on the coat.'
7	Mat-first	null	Il papà saluta il figlio, mentre ø va in bicicletta.
		overt	Il papà saluta il figlio, mentre lui va in bicicletta.
			'The father greets the son, while he rides the bike.'
8	Mat-first	null	Il poliziotto vede il ladro, mentre ø corre.
		overt	Il poliziotto vede il ladro, mentre lui corre.
			'The policeman sees the thief, while he runs.'
9	Mat-first	null	La maestra indica l'alunna, mentre ø parla.
		overt	La maestra indica l'alunna, mentre lei parla.
			'The teacher points at the pupil, while she talks.'
10	Mat-first	null	L'allenatore sta parlando all'atleta, mentre ø sta bevendo.
		overt	L'allenatore sta parlando all'atleta, mentre lui sta bevendo.
			'The trainer is talking to the athlete, while he is drinking.'
11	Emb-first	null	Appena ø chiude la borsa, il fattorino dà il denaro al cassiere.
		overt	Appena lui chiude la borsa, il fattorino dà il denaro al cassiere.
			'As soon as he closes the bag, the errand boy gives the money to the cashier.'
12	Emb-first	null	Mentre ø apre la porta, il portiere saluta il postino.
		overt	Mentre lui apre la porta, il portiere saluta il postino.
			'While he opens the door, the doorman greets the postman.'

continued on the next page

Num.	Direction	Pronoun	Sentence
13	Emb-first	null	Mentre ø guarda l'orologio, l'anziana signora si avvicina alla donna delle pulizie.
		overt	Mentre lei guarda l'orologio, l'anziana signora si avvicina alla donna delle pulizie
			'While she looks at the watch, the elderly woman moves toward the cleaning lady.'
14	Emb-first	null	Appena ø gira l'angolo, il poliziotto vede il ladro.
		overt	Appena lui gira l'angolo, il poliziotto vede il ladro.
			'As soon as he turns the corner, the policeman sees the theif.'
15	Emb-first	null	Mentre ø versa il vino nel bicchiere, il cliente paga il conto al cameriere.
		overt	Mentre lui versa il vino nel bicchiere, il cliente paga il conto al cameriere.
			'While he pours the wine in the glass, the client pays the waiter the bill.'
16*	Emb-first	null	Mentre ø sabdiglia, il controllore prende il biglietto dal passeggero.
		overt	Mentre lui sabdiglia, il controllore prende il biglietto dal passeggero.
			'While he yawns, the inspector takes the ticket from the passenger.'
17	Emb-first	null	Mentre ø esce dall'ascensore, l'infermiera urta la donna delle pulizie.
		overt	Mentre lei esce dall'ascensore, l'infermiera urta la donna delle pulizie.
			'While she leaves the elevator, the nurse bumps into the cleaning lady.'
18*	Emb-first	null	Mentre ø cammina nel parco, il guardiano vede il barbone.
		overt	Mentre lui cammina nel parco, il guardiano vede il barbone.
			'While he walks in the park, the watchman sees the homeless man.'
19	Emb-first	null	Mentre ø aspetta l'autobus, il prete dice una cosa al turista.
		overt	Mentre lui aspetta l'autobus, il prete dice una cosa al turista.
			'While he waits for the bus, the priest says something to the turist.'
20	Emb-first	null	Appena ø apre la porta, il papà augura buon compleanno al figlio.
		overt	Appena lui apre la porta, il papà augura buon compleanno al figlio.
			'As soon as he opens the door, the father wishes the son happy birthday.'

Table B.2: Pronominal items adapted from Tsimpli et al.'s (2004) sentence interpretation task with added translations. Asterisks are used to mark the two items that were not included in statistical analyses due to coding errors.

C Locally ambiguous stimuli

Num.	Predicate	Attach.	Sentence
1	non-perceptual	low	Gianni vive con il collega della biologa che correva sporca di fango.
		high	Gianni vive con il collega della biologa che correva sporco di fango.
	perceptual	low	Gianni ha visto il collega della biologa che correva sporca di fango.
		high	Gianni ha visto il collega della biologa che correva sporco di fango.
			'Gianni (lives with / saw) the collegue of the biologist that was running covered in mud.'
2	non-perceptual	low	Simona lavora con la figlia del deputato che parlava emozionato al telefono.
		high	Simona lavora con la figlia del deputato che parlava emozionata al telefono.
	perceptual	low	Simona ha sentito la figlia del deputato che parlava emozionato al telefono.
		high	Simona ha sentito la figlia del deputato che parlava emozionata al telefono.
			'Simona (works with / heard) the daughter of the MP that was talking excitedly on the phone.'
3	non-perceptual	low	Pietro studia con il figlio della segretaria che mangiava seduta sulla panchina.
		high	Pietro studia con il figlio della segretaria che mangiava seduto sulla panchina.
	perceptual	low	Pietro guardava il figlio della segretaria che mangiava seduta sulla panchina.
		high	Pietro guardava il figlio della segretaria che mangiava seduto sulla panchina.
			'Pietro (studies with / watched) the son of the secretary that was eating sitting on the bench.'
4	non-perceptual	low	Silvio frequenta la cugina dell'infermiere che fumava affacciato alla finestra.
		high	Silvio frequenta la cugina dell'infermiere che fumava affacciata alla finestra.
	perceptual	low	Silvio ha beccato la cugina dell'infermiere che fumava affacciato alla finestra.
		high	Silvio ha beccato la cugina dell'infermiere che fumava affacciata alla finestra.
			'Silvio (is dating / caught) the cousin of the nurse that was smoking facing the window. '
5	non-perceptual	low	Luigi lavora per il nonno della ballerina che sedeva rilassata sul balcone.
		high	Luigi lavora per il nonno della ballerina che sedeva rilassato sul balcone.
	perceptual	low	Luigi ha ritratto il nonno della ballerina che sedeva rilassata sul balcone.
		high	Luigi ha ritratto il nonno della ballerina che sedeva rilassato sul balcone.
			'Luigi (works for / portrayed) the grandfather of the dancer that was sitting relaxed on the balcony.'
6	non-perceptual	low	Maria convive con la nipote del bottegaio che urlava ubriaco nella piazza.
		high	Maria convive con la nipote del bottegaio che urlava ubriaca nella piazza.
	perceptual	low	Maria ha registrato la nipote del bottegaio che urlava ubriaco nella piazza.
		high	Maria ha registrato la nipote del bottegaio che urlava ubriaca nella piazza.
			'Maria (lives with / recorded) the niece of the shopkeeper that was shouting drunk in the piazza.'
7	non-perceptual	low	Giulia frequenta il vicino della signora che studiava circondata di quaderni.
		high	Giulia frequenta il vicino della signora che studiava circondato di quaderni.

continued on the next page

Num.	Predicate	Attach.	Sentence
	perceptual	low	Giulia guardava il vicino della signora che studiava circondata di quaderni.
		high	Giulia guardava il vicino della signora che studiava circondato di quaderni.
			'Giulia (is dating / watched) the neighbour of the woman that was studying surrounded by notebooks'
8	non-perceptual	low	Claudio ha sposato la figliastra del sarto che fischiava contento nel suo giardino.
		high	Claudio ha sposato la figliastra del sarto che fischiava contenta nel suo giardino.
	perceptual	low	Claudio ha ripreso la figliastra del sarto che fischiava contento nel suo giardino.
		high	Claudio ha ripreso la figliastra del sarto che fischiava contenta nel suo giardino.
			'Claudio (married / filmed) the stepdaughter of the tailor that was whistling contentedly in their garden.'
9	non-perceptual	low	Giorgina esce con il maestro della bambina che camminava distratta dal cellulare.
		high	Giorgina esce con il maestro della bambina che camminava distratto dal cellulare.
	perceptual	low	Giorgina osservava il maestro della bambina che camminava distratta dal cellulare.
		high	Giorgina osservava il maestro della bambina che camminava distratto dal cellulare.
			'Giorgina (goes out with / observed) the teacher of the girl that was walking distracted by their phone'
10	non-perceptual	low	Martina conosce l'alunna del musicista che girava silenzioso per i corridoi.
		high	Martina conosce l'alunna del musicista che girava silenziosa per i corridoi.
	perceptual	low	Martina ha visto l'alunna del musicista che girava silenzioso per i corridoi.
		high	Martina ha visto l'alunna del musicista che girava silenziosa per i corridoi.
			'Martina (knows / saw) the pupil of the musician that was walking silently around the halls.'
11	non-perceptual	low	Davide collabora con l'amico della giornalista che pregava seduta nella poltrona.
		high	Davide collabora con l'amico della giornalista che pregava seduto nella poltrona.
	perceptual	low	Davide ha disegnato l'amico della giornalista che pregava seduta nella poltrona.
		high	Davide ha disegnato l'amico della giornalista che pregava seduto nella poltrona.
			'Davide (works with / drew) the friend of the journalist that was praying sitting in the armchair.'
12	non-perceptual	low	Aurora studia con la nipote del commesso che cantava sereno davanti alla folla.
		high	Aurora studia con la nipote del commesso che cantava serena davanti alla folla.
	perceptual	low	Aurora ascoltava la nipote del commesso che cantava sereno davanti alla folla.
		high	Aurora ascoltava la nipote del commesso che cantava serena davanti alla folla.
			'Aurora (studies with / listened to) the niece of the clerk that was singing serenely in front of the crowd.'
13	non-perceptual	low	Matteo vive con il cognato della casalinga che guidava impazzita in tangenziale.
		high	Matteo vive con il cognato della casalinga che guidava impazzito in tangenziale.
	perceptual	low	Matteo ha visto il cognato della casalinga che guidava impazzita in tangenziale.
		high	Matteo ha visto il cognato della casalinga che guidava impazzito in tangenziale.
			'Matteo (lives with / saw) the brother-in-law of the housewife that was driving crazy on the bypass.'

continued on the next page

Num.	Predicate	Attach.	Sentence
14	non-perceptual	low	Emilio è sposato con la cugina del ragazzo che fischiava nervoso nel cortile.
		high	Emilio è sposato con la cugina del ragazzo che fischiava nervosa nel cortile.
	perceptual	low	Emilio ha registrato la cugina del ragazzo che fischiava nervoso nel cortile.
		high	Emilio ha registrato la cugina del ragazzo che fischiava nervosa nel cortile.
			'Emilio (is married to / recorded) the cousin of the boy that was whistling nervously in the courtyard.'
15	non-perceptual	low	Filippo lavora per il nonno della ginnasta che leggeva sdraiata sotto l'albero.
		high	Filippo lavora per il nonno della ginnasta che leggeva sdraiato sotto l'albero.
	perceptual	low	Filippo ha ritratto il nonno della ginnasta che leggeva sdraiata sotto l'albero.
		high	Filippo ha ritratto il nonno della ginnasta che leggeva sdraiato sotto l'albero.
			'Filippo (works for / potrayed) the grandfather of the gymnast that was reading lying under the tree.'
16	non-perceptual	low	Chiara vive con la nipote del signore che correva fradicio di sudore.
		high	Chiara vive con la nipote del signore che correva fradicia di sudore.
	perceptual	low	Chiara immaginava la nipote del signore che correva fradicio di sudore.
		high	Chiara immaginava la nipote del signore che correva fradicia di sudore.
			'Chiara (lives with / imagined) the niece of the man that was running drenched in sweat.'
17*	non-perceptual	low	Carla ha sposato il collega della postina che cantava seduta tra i suoi amici.
		high	Carla ha sposato il collega della postina che cantava seduto tra i suoi amici.
	perceptual	low	Carla ha ripreso il collega della postina che cantava seduta tra i suoi amici.
		high	Carla ha ripreso il collega della postina che cantava seduto tra i suoi amici.
			'Carla (married / filmed) the colleague of the postwoman that was singing sitting among her friends.'
18*	non-perceptual	low	Claudio ha sposato la cognata dell'infermiere che vagava confuso nel bosco.
		high	Claudio ha sposato la cognata dell'infermiere che vagava confusa nel bosco.
	perceptual	low	Claudio ha sognato la cognata dell'infermiere che vagava confuso nel bosco.
		high	Claudio ha sognato la cognata dell'infermiere che vagava confusa nel bosco.
			'Claudio (married / dreamed of) the sister-in-law of the nurse that was wandering confused in the woods.'
19	non-perceptual	low	Roberta esce con il maestro della bambina che scriveva appoggiata alla lavagna.
		high	Roberta esce con il maestro della bambina che scriveva appoggiato alla lavagna.
	perceptual	low	Roberta osservava il maestro della bambina che scriveva appoggiata alla lavagna.
		high	Roberta osservava il maestro della bambina che scriveva appoggiato alla lavagna.
			'Roberta (goes out with / observed) the teacher of the girl that was writing leaning against the blackboard'
20	non-perceptual	low	Daniele frequenta la cugina del cameriere che dormiva tranquillo nel giardino.
		high	Daniele frequenta la cugina del cameriere che dormiva tranquilla nel giardino.
	perceptual	low	Daniele ha beccato la cugina del cameriere che dormiva tranquillo nel giardino.
		high	Daniele ha beccato la cugina del cameriere che dormiva tranquilla nel giardino.

continued on the next page

Num.	Predicate	Attach.	Sentence
			'Daniele (is dating / caught) the cousin of the waiter that was sleeping peacefully in the garden.'
21	non-perceptual	low	Gianni lavora con lo zio della segretaria che balbettava nervosa davanti alla polizia.
		high	Gianni lavora con lo zio della segretaria che balbettava nervoso davanti alla polizia.
	perceptual	low	Gianni ascoltava lo zio della segretaria che balbettava nervosa davanti alla polizia.
		high	Gianni ascoltava lo zio della segretaria che balbettava nervoso davanti alla polizia.
			'Gianni (works with / listened to) the uncle of the secretary that was babbling nervously in front of the police.'
22	non-perceptual	low	Elisa convive con la figlia del postino che parlava circondato dai giornalisti.
		high	Elisa convive con la figlia del postino che parlava circondata dai giornalisti.
	perceptual	low	Elisa ha filmato la figlia del postino che parlava circondato dai giornalisti.
		high	Elisa ha filmato la figlia del postino che parlava circondata dai giornalisti.
			'Elisa (lives with / filmed) the daughter of the postman that was talking surrounded by the journalists.'
23	non-perceptual	low	Alessia è fidanzata con il cugino della signora che camminava agitata lungo la strada.
		high	Alessia è fidanzata con il cugino della signora che camminava agitato lungo la strada.
	perceptual	low	Alessia ha intravisto il cugino della signora che camminava agitata lungo la strada.
		high	Alessia ha intravisto il cugino della signora che camminava agitato lungo la strada.
			'Alessia (is in a relationship with / glimpsed) the cousin of the woman that was walking agitatedly along the road.'
24	non-perceptual	low	Viola studia con la figliastra del cuoco che balbettava arrabbiato al cellulare.
		high	Viola studia con la figliastra del cuoco che balbettava arrabbiata al cellulare.
	perceptual	low	Viola ascoltava la figliastra del cuoco che balbettava arrabbiato al cellulare.
		high	Viola ascoltava la figliastra del cuoco che balbettava arrabbiata al cellulare.
			'Viola (studies with / listened to) the stepdaughter of the cook that was babbling angrily on the cellphone.'
25	non-perceptual	low	Giorgio conosce l'amico della scienziata che giaceva pallida sul pavimento.
		high	Giorgio conosce l'amico della scienziata che giaceva pallido sul pavimento.
	perceptual	low	Giorgio ha trovato l'amico della scienziata che giaceva pallida sul pavimento.
		high	Giorgio ha trovato l'amico della scienziata che giaceva pallido sul pavimento.
			'Giorgio (knows / found) the friend of the scientist that was lying pale on the floor.'
26*	non-perceptual	low	Silvio ha sposato la collega del commesso che passeggiava inquieto nel bosco.
		high	Silvio ha sposato la collega del commesso che passeggiava inquieta nel bosco.
	perceptual	low	Silvio ha sognato la collega del commesso che passeggiava inquieto nel bosco.
		high	Silvio ha sognato la collega del commesso che passeggiava inquieta nel bosco.
			'Silvia (married / dreamed of) the colleague of the clerk that was walking restlessly in the woods.'
27	non-perceptual	low	Daniele lavora con l'amico della pensionata che urlava confusa per l'ospedale.
		high	Daniele lavora con l'amico della pensionata che urlava confuso per l'ospedale.
	perceptual	low	Daniele ascoltava l'amico della pensionata che urlava confusa per l'ospedale.
		high	Daniele ascoltava l'amico della pensionata che urlava confuso per l'ospedale.

continued on the next page

Num.	Predicate	Attach.	Sentence
			'Daniele (works for / listened to) the friend of the pensioner that was shouting confusedly about the hospital.'
28	non-perceptual	low	Maria si allena con la figlia del farmacista che aspettava ansioso nell'ambulatorio.
		high	Maria si allena con la figlia del farmacista che aspettava ansiosa nell'ambulatorio.
	perceptual	low	Maria ha sorpreso la figlia del farmacista che aspettava ansioso nell'ambulatorio.
		high	Maria ha sorpreso la figlia del farmacista che aspettava ansiosa nell'ambulatorio.
			'Maria (works out with / surprised) the daughter of the pharmacist that was waiting anxiously in the waiting room.'
29	non-perceptual	low	Giulia è fidanzata con il vicino della bottegaia che girava vestita solo in pigiama.
		high	Giulia è fidanzata con il vicino della bottegaia che girava vestito solo in pigiama.
	perceptual	low	Giulia ha intravisto il vicino della bottegaia che girava vestita solo in pigiama.
		high	Giulia ha intravisto il vicino della bottegaia che girava vestito solo in pigiama.
			'Giulia (is in a relationship with / glimpsed) the neighbour of the shopkeeper that was walking around dressed in just their pajamas.'
30	non-perceptual	low	Camilla si allena con la figliastra del vedovo che girava scalzo nell'ufficio.
		high	Camilla si allena con la figliastra del vedovo che girava scalza nell'ufficio.
	perceptual	low	Camilla ha sorpreso la figliastra del vedovo che girava scalzo nell'ufficio.
		high	Camilla ha sorpreso la figliastra del vedovo che girava scalza nell'ufficio.
			'Camilla (works out with / surprised) the stepdaughter of the widower that was walking around the office barefoot.'
31	non-perceptual	low	Federico conosce l'amico della musicista che studiava assonnata nella biblioteca.
		high	Federico conosce l'amico della musicista che studiava assonnato nella biblioteca.
	perceptual	low	Federico ha trovato l'amico della musicista che studiava assonnata nella biblioteca.
		high	Federico ha trovato l'amico della musicista che studiava assonnato nella biblioteca.
			'Federico (knows / found) the friend of the musician that was studying sleepily in the library.'
32	non-perceptual	low	Aurora si esercita con la vicina dello psicologo che vagava scalzo nel giardino.
		high	Aurora si esercita con la vicina dello psicologo che vagava scalza nel giardino.
	perceptual	low	Aurora ha immaginato la vicina dello psicologo che vagava scalzo nel giardino.
		high	Aurora ha immaginato la vicina dello psicologo che vagava scalza nel giardino.
			'Aurora (works out with / imagined) the neighbour of the psychologist that was wondering barefoot in the garden.'

Table C.1: Critical PR/RC items. An asterisk is used to indicate the three items which were lost for 12 control participants due to a coding error.

Number	Pronoun	Antecedent	Sentence
1	null	non-SpecIP	Bruno ha incontrato Beatrice mentre tornava brilla dopo una festa.
		SpecIP	Bruno ha incontrato Beatrice mentre tornava brillo dopo una festa.
	overt	non-SpecIP	Bruno ha incontrato Beatrice mentre lei tornava brilla dopo una festa.
		SpecIP	Bruno ha incontrato Beatrice mentre lui tornava brillo dopo una festa.
			'Bruno bumped into Beatrice while s/he were coming back tipsy from a party.'
2	null	non-SpecIP	Clarissa ha chiamato Fabio mentre usciva nervoso dal colloquio.
		SpecIP	Clarissa ha chiamato Fabio mentre usciva nervosa dal colloquio.
	overt	non-SpecIP	Clarissa ha chiamato Fabio mentre lui usciva nervoso dal colloquio.
		SpecIP	Clarissa ha chiamato Fabio mentre lei usciva nervosa dal colloquio.
			'Clarissa called Fabio as s/he was nervously walking out of the interview.'
3	null	non-SpecIP	Mauro ha visto Costanza mentre rientrava stanca dopo il lavoro.
		SpecIP	Mauro ha visto Costanza mentre rientrava stanco dopo il lavoro.
	overt	non-SpecIP	Mauro ha visto Costanza mentre lei rientrava stanca dopo il lavoro.
		SpecIP	Mauro ha visto Costanza mentre lui rientrava stanco dopo il lavoro.
			'Mauro saw Coztanza while s/he was coming back tired after work.'
4	null	non-SpecIP	Cristina ha telefonato ad Angelo mentre girava tranquillo per il parco.
		SpecIP	Cristina ha telefonato ad Angelo mentre girava tranquilla per il parco.
	overt	non-SpecIP	Cristina ha telefonato ad Angelo mentre lui girava tranquillo per il parco.
		SpecIP	Cristina ha telefonato ad Angelo mentre lei girava tranquilla per il parco.
			'Cristina called Angelo while s/he was walking peacfully in the park.'
5	null	non-SpecIP	Iacopo ha salutato Giovanna mentre rientrava sudata dopo la palestra.
		SpecIP	Iacopo ha salutato Giovanna mentre rientrava sudato dopo la palestra.
	overt	non-SpecIP	Iacopo ha salutato Giovanna mentre lei rientrava sudata dopo la palestra.
		SpecIP	Iacopo ha salutato Giovanna mentre lui rientrava sudato dopo la palestra.
			'Iacopo greeted Giovanna while s/he was coming back sweaty after the gym.'
6	null	non-SpecIP	Stefania ha minacciato Renato mentre mangiava seduto al tavolo.
		SpecIP	Stefania ha minacciato Renato mentre mangiava seduta al tavolo.
	overt	non-SpecIP	Stefania ha minacciato Renato mentre lui mangiava seduto al tavolo.
		SpecIP	Stefania ha minacciato Renato mentre lei mangiava seduta al tavolo.
			'Stefania threatened Renato while s/he was eating sat at the table.'
7	null	non-SpecIP	Sergio ha beccato Virginia mentre passeggiava silenziosa nel parco.
		SpecIP	Sergio ha beccato Virginia mentre passeggiava silenzioso nel parco.
	overt	non-SpecIP	Sergio ha beccato Virginia mentre lei passeggiava silenziosa nel parco.
		SpecIP	Sergio ha beccato Virginia mentre lui passeggiava silenzioso nel parco.
			'Sergio surprised Virginia while s/he was walking silently in the park.'
8	null	non-SpecIP	Vittoria parlava con Simone mentre giaceva comodo sul divano.
		SpecIP	Vittoria parlava con Simone mentre giaceva comoda sul divano.
	overt	non-SpecIP	Vittoria parlava con Simone mentre lui giaceva comodo sul divano.
		SpecIP	Vittoria parlava con Simone mentre lei giaceva comoda sul divano.
			'Vittoria was talking with Simone while s/he was lying comfortably on the couch.'

continued on the next page

Number	Pronoun	Antecedent	Sentence
9	null	non-SpecIP	Adriano fissava Arianna mentre fumava agitata davanti a un bar.
		SpecIP	Adriano fissava Arianna mentre fumava agitato davanti a un bar.
	overt	non-SpecIP	Adriano fissava Arianna mentre fumava lei agitata davanti a un bar.
		SpecIP	Adriano fissava Arianna mentre lui fumava agitato davanti a un bar.
			'Adriano stared at Arianna while s/he was shaken and smoking in front of a bar.'
10	null	non-SpecIP	Iolanda ha baciato Michele mentre leggeva disteso sul divano.
		SpecIP	Iolanda ha baciato Michele mentre leggeva distesa sul divano
	overt	non-SpecIP	Iolanda ha baciato Michele mentre lui leggeva disteso sul divano
		SpecIP	Iolanda ha baciato Michele mentre lei leggeva distesa sul divano
			'Iolanda kissed Michele while s/he was reading lying on the couch.'
11	null	non-SpecIP	Stefano ha abbracciato Rachele mentre piangeva affranta davanti all'ospedale.
		SpecIP	Stefano ha abbracciato Rachele mentre piangeva affranto davanti all'ospedale.
	overt	non-SpecIP	Stefano ha abbracciato Rachele mentre lei piangeva affranta davanti all'ospedale.
		SpecIP	Stefano ha abbracciato Rachele mentre lui piangeva affranto davanti all'ospedale.
			'Stefano huggged Rachele while s/he was distraught and crying in front of the hospital.'
12	null	non-SpecIP	Valeria ha sorriso a Tommaso mentre si rilassava disteso vicino alla piscina.
		SpecIP	Valeria ha sorriso a Tommaso mentre si rilassava distesa vicino alla piscina.
	overt	non-SpecIP	Valeria ha sorriso a Tommaso mentre lui si rilassava disteso vicino alla piscina.
		SpecIP	Valeria ha sorriso a Tommaso mentre lei si rilassava distesa vicino alla piscina.
			'Valerai smiled at Tommaso while s/he was relaxing and lying down by the pool.'
13	null	non-SpecIP	Domenico ha fotografato Fatima mentre fumava seduta sul balcone.
		SpecIP	Domenico ha fotografato Fatima mentre fumava seduto sul balcone.
	overt	non-SpecIP	Domenico ha fotografato Fatima mentre lei fumava seduta sul balcone.
		SpecIP	Domenico ha fotografato Fatima mentre lui fumava seduto sul balcone.
			'Domenico photographed Fatima while s/he was smoking sat on the balcony.'
14	null	non-SpecIP	Ilaria ha salutato Emanuele mentre usciva vestito per la palestra.
		SpecIP	Ilaria ha salutato Emanuele mentre usciva vestita per la palestra.
	overt	non-SpecIP	Ilaria ha salutato Emanuele mentre lui usciva vestito per la palestra.
		SpecIP	Ilaria ha salutato Emanuele mentre lei usciva vestita per la palestra.
			'Ilaria greeted Emanuele while s/he was leaving dressed for the gym.'
15	null	non-SpecIP	Fabrizio guardava Teresa mentre rabbrividiva avvolta in una coperta.
		SpecIP	Fabrizio guardava Teresa mentre rabbrividiva avvolto in una coperta.
	overt	non-SpecIP	Fabrizio guardava Teresa mentre lei rabbrividiva avvolta in una coperta.
		SpecIP	Fabrizio guardava Teresa mentre lui rabbrividiva avvolto in una coperta.
			'Fabrizio watched Teresa while s/he shivered wrapped up in a blanket.'
16	null	non-SpecIP	Agata ha chiamato Giovanni mentre si rilassava sdraiato in spiaggia.
		SpecIP	Agata ha chiamato Giovanni mentre si rilassava sdraiata in spiaggia.
	overt	non-SpecIP	Agata ha chiamato Giovanni mentre lui si rilassava sdraiato in spiaggia.
		SpecIP	Agata ha chiamato Giovanni mentre lei si rilassava sdraiata in spiaggia.
			'Agata called Giovanni while s/he was relaxing sprawled on the beach.'

continued on the next page

Number	Pronoun	Antecedent	Sentence
17	null	non-SpecIP	Giuseppe camminava con Elena mentre parlava emozionata della sua famiglia.
		SpecIP	Giuseppe camminava con Elena mentre parlava emozionato della sua famiglia.
	overt	non-SpecIP	Giuseppe camminava con Elena mentre lei parlava emozionata della sua famiglia.
		SpecIP	Giuseppe camminava con Elena mentre lui parlava emozionato della sua famiglia.
			'Giuseppe was walking with Elena while s/he spoke emotionally about his/her family.'
18	null	non-SpecIP	Irene ha visto Riccardo mentre guidava impazzito in autostrada.
		SpecIP	Irene ha visto Riccardo mentre guidava impazzita in autostrada.
	overt	non-SpecIP	Irene ha visto Riccardo mentre lui guidava impazzito in autostrada.
		SpecIP	Irene ha visto Riccardo mentre lei guidava impazzita in autostrada.
			'Irene saw Riccarod while s/he was driving crazily on the highway'
19	null	non-SpecIP	Bruno ha inviato una e-mail a Vittoria mentre lavorava chiusa in ufficio.
		SpecIP	Bruno ha inviato una e-mail a Vittoria mentre lavorava chiuso in ufficio.
	overt	non-SpecIP	Bruno ha inviato una e-mail a Vittoria mentre lei lavorava chiusa in ufficio.
		SpecIP	Bruno ha inviato una e-mail a Vittoria mentre lui lavorava chiuso in ufficio.
			'Bruno sent an email to Vittoria while s/he was working closed in the office.'
20	null	non-SpecIP	Virginia ha fatto un cenno a Dario mentre saliva nervoso sul palcoscenico.
		SpecIP	Virginia ha fatto un cenno a Dario mentre saliva nervosa sul palcoscenico.
	overt	non-SpecIP	Virginia ha fatto un cenno a Dario mentre lui saliva nervoso sul palcoscenico.
		SpecIP	Virginia ha fatto un cenno a Dario mentre lei saliva nervosa sul palcoscenico.
			'Virginia nodded to Dario while s/he was walking nervously on stage.'
21	null	non-SpecIP	Angelo guardava Eleonora mentre pranzava tranquilla sulla terrazza.
		SpecIP	Angelo guardava Eleonora mentre pranzava tranquillo sulla terrazza.
	overt	non-SpecIP	Angelo guardava Eleonora mentre lei pranzava tranquilla sulla terrazza.
		SpecIP	Angelo guardava Eleonora mentre lui pranzava tranquillo sulla terrazza.
			'Angelo watched Eleonora while s/he was taking lunch peacfully on the terrace.'
22	null	non-SpecIP	Cristina fissava Enrico mentre aspettava seduto fuori dal tribunale.
		SpecIP	Cristina fissava Enrico mentre aspettava seduta fuori dal tribunale.
	overt	non-SpecIP	Cristina fissava Enrico mentre lui aspettava seduto fuori dal tribunale.
		SpecIP	Cristina fissava Enrico mentre lei aspettava seduta fuori dal tribunale.
			'Cristina stared at Enrico while s/he waited sat outside the courthouse.'
23	null	non-SpecIP	Renato ha mandato un sms a Clarissa mentre camminava affamata verso il ristorante.
		SpecIP	Renato ha mandato un sms a Clarissa mentre camminava affamato verso il ristorante.
	overt	non-SpecIP	Renato ha mandato un sms a Clarissa mentre lei camminava affamata verso il ristorante.
		SpecIP	Renato ha mandato un sms a Clarissa mentre lui camminava affamato verso il ristorante.
			'Renato sent a text message to Clarissa while s/he was walking hungrily toward the restaurant.'
24	null	non-SpecIP	Valeria ha salutato Adriano mentre tornava affamato dalla palestra.
		SpecIP	Valeria ha salutato Adriano mentre tornava affamata dalla palestra.
	overt	non-SpecIP	Valeria ha salutato Adriano mentre lui tornava affamato dalla palestra.

continued on the next page

Number	Pronoun	Antecedent	Sentence
		SpecIP	Valeria ha salutato Adriano mentre lei tornava affamata dalla palestra.
			'Valeria greeted Adriano while s/he was coming back hungry from the gym.'
25	null	non-SpecIP	Alberto ha sorriso a Rachele mentre parlava circondata dai suoi amici.
		SpecIP	Alberto ha sorriso a Rachele mentre parlava circondato dai suoi amici.
	overt	non-SpecIP	Alberto ha sorriso a Rachele mentre lei parlava circondata dai suoi amici.
		SpecIP	Alberto ha sorriso a Rachele mentre lui parlava circondato dai suoi amici.
			'Alberto smiled at Rachele while s/he was talking surrounded by her/his friends.'
26	null	non-SpecIP	Iolanda ha salutato Michele mentre rientrava rilassato dopo le vacanze.
		SpecIP	Iolanda ha salutato Michele mentre rientrava rilassata dopo le vacanze.
	overt	non-SpecIP	Iolanda ha salutato Michele mentre lui rientrava rilassato dopo le vacanze.
		SpecIP	Iolanda ha salutato Michele mentre lei rientrava rilassata dopo le vacanze.
			'Iolanda greeted Michele while s/he was returning relaxed after the holiday.'
27	null	non-SpecIP	Tommaso ha beccato Arianna mentre correva scalza in spiaggia.
		SpecIP	Tommaso ha beccato Arianna mentre correva scalzo in spiaggia.
	overt	non-SpecIP	Tommaso ha beccato Arianna mentre lei correva scalza in spiaggia.
		SpecIP	Tommaso ha beccato Arianna mentre lui correva scalzo in spiaggia.
			'Tommaso surprised arianna while s/he was running barefoot on the beach.'
28*	null	non-SpecIP	Teresa ha fatto l'occhiolino a Domenico mentre passeggiava rilassato in centro.
		SpecIP	Teresa ha fatto l'occhiolino a Domenico mentre passeggiava rilassata in centro.
	overt	non-SpecIP	Teresa ha fatto l'occhiolino a Domenico mentre lui passeggiava rilassato in centro.
		SpecIP	Teresa ha fatto l'occhiolino a Domenico mentre lei passeggiava rilassata in centro.
			'Teresa winked at Domenico while s/he was strolling relaxed in the city center.'
29	null	non-SpecIP	Fabrizio ha visto Fatima mentre aspettava ansiosa nella sala di attesa.
		SpecIP	Fabrizio ha visto Fatima mentre aspettava ansioso nella sala di attesa.
	overt	non-SpecIP	Fabrizio ha visto Fatima mentre lei aspettava ansiosa nella sala di attesa.
		SpecIP	Fabrizio ha visto Fatima mentre lui aspettava ansioso nella sala di attesa.
			'Fabrizio saw Fatima while s/he was waiting anxiously in the waiting room.'
30	null	non-SpecIP	Elena ha baciato Riccardo mentre si vestiva assonnato prima del lavoro.
		SpecIP	Elena ha baciato Riccardo mentre si vestiva assonnata prima del lavoro.
	overt	non-SpecIP	Elena ha baciato Riccardo mentre lui si vestiva assonnato prima del lavoro.
		SpecIP	Elena ha baciato Riccardo mentre lei si vestiva assonnata prima del lavoro.
			'Elena kissed Riccardo while s/he was sleeply getting dressed before work.'
31	null	non-SpecIP	Giuseppe ha mandato un sms a Irene mentre mangiava appoggiata al bancone del bar.
		SpecIP	Giuseppe ha mandato un sms a Irene mentre mangiava appoggiato al bancone del bar.
	overt	non-SpecIP	Giuseppe ha mandato un sms a Irene mentre lei mangiava appoggiata al bancone del bar.
		SpecIP	Giuseppe ha mandato un sms a Irene mentre lui mangiava appoggiato al bancone del bar.
			'Giuseppe sent a text message to Irene while s/he was eating leaning against the bar

continued on the next page

Number	Pronoun	Antecedent	Sentence
			counter.'
32	null	non-SpecIP	Adele ha incontrato Leonardo mentre girava sereno nel parco.
		SpecIP	Adele ha incontrato Leonardo mentre girava serena nel parco.
	overt	non-SpecIP	Adele ha incontrato Leonardo mentre lui girava sereno nel parco.
		SpecIP	Adele ha incontrato Leonardo mentre lei girava serena nel parco.
			'Adele bumped into Leonardo while s/he was strolling serenely around the park.'

Table C.2: Critical Pronominal items. An asterisk is used to indicate the item which was not included in the statistical analysis due to a coding error.

D Supplementary tables and figures

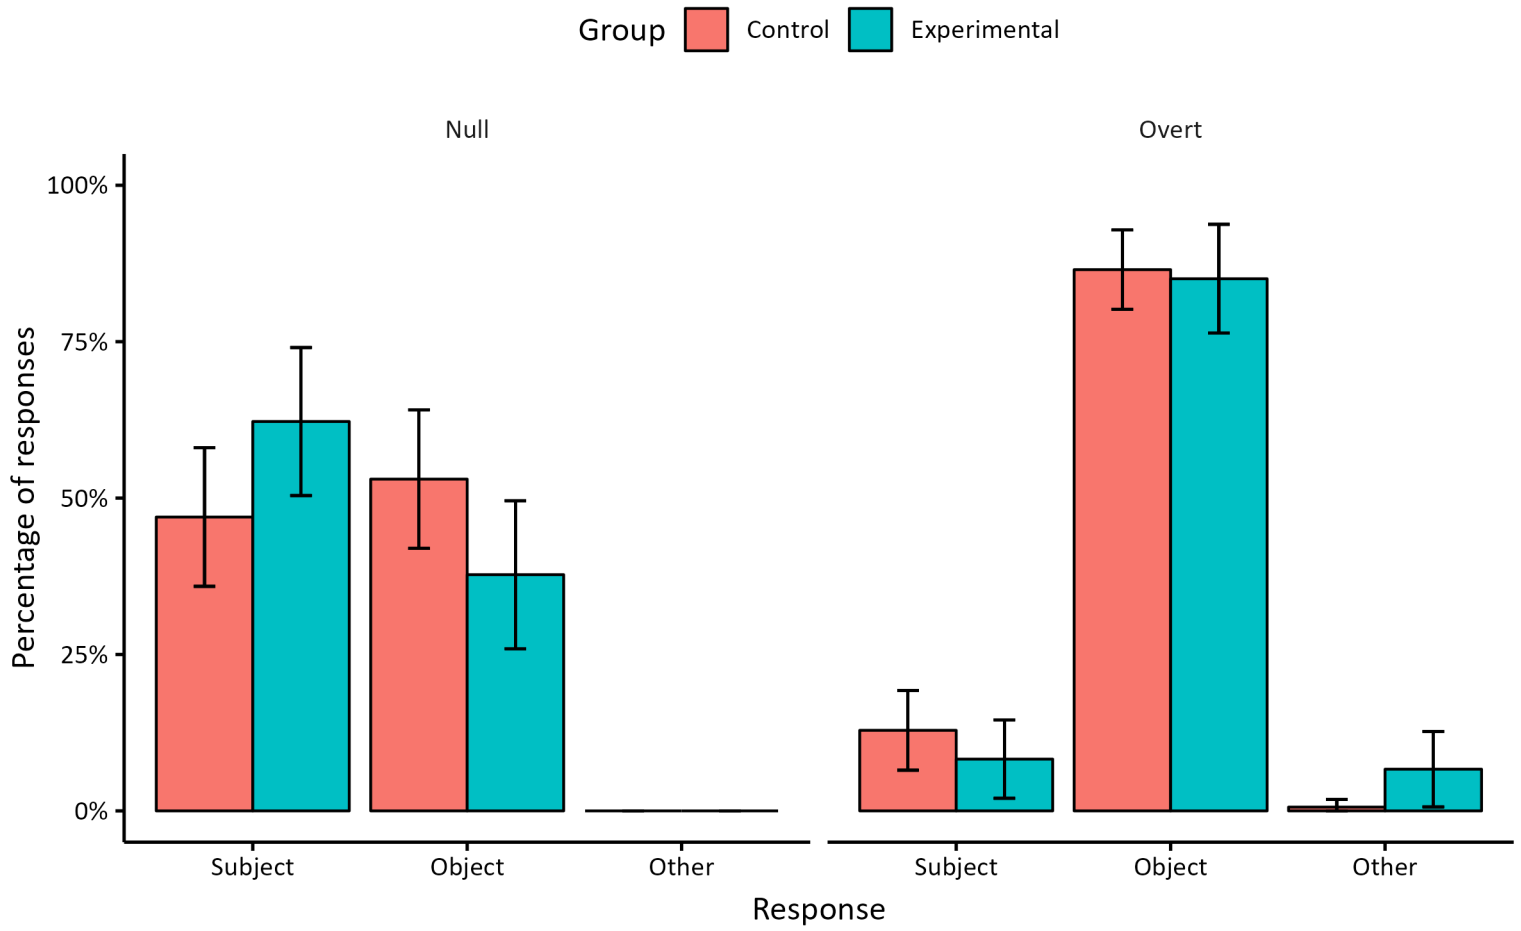

Figure D.1: Responses to globally ambiguous matrix-first pronominal items in chapter 3 with the [+TS] readings separated into 'object' and 'other'. Please note the subject values trivially match those reported in Figure 3.1a and 'object' responses include both direct and indirect objects.

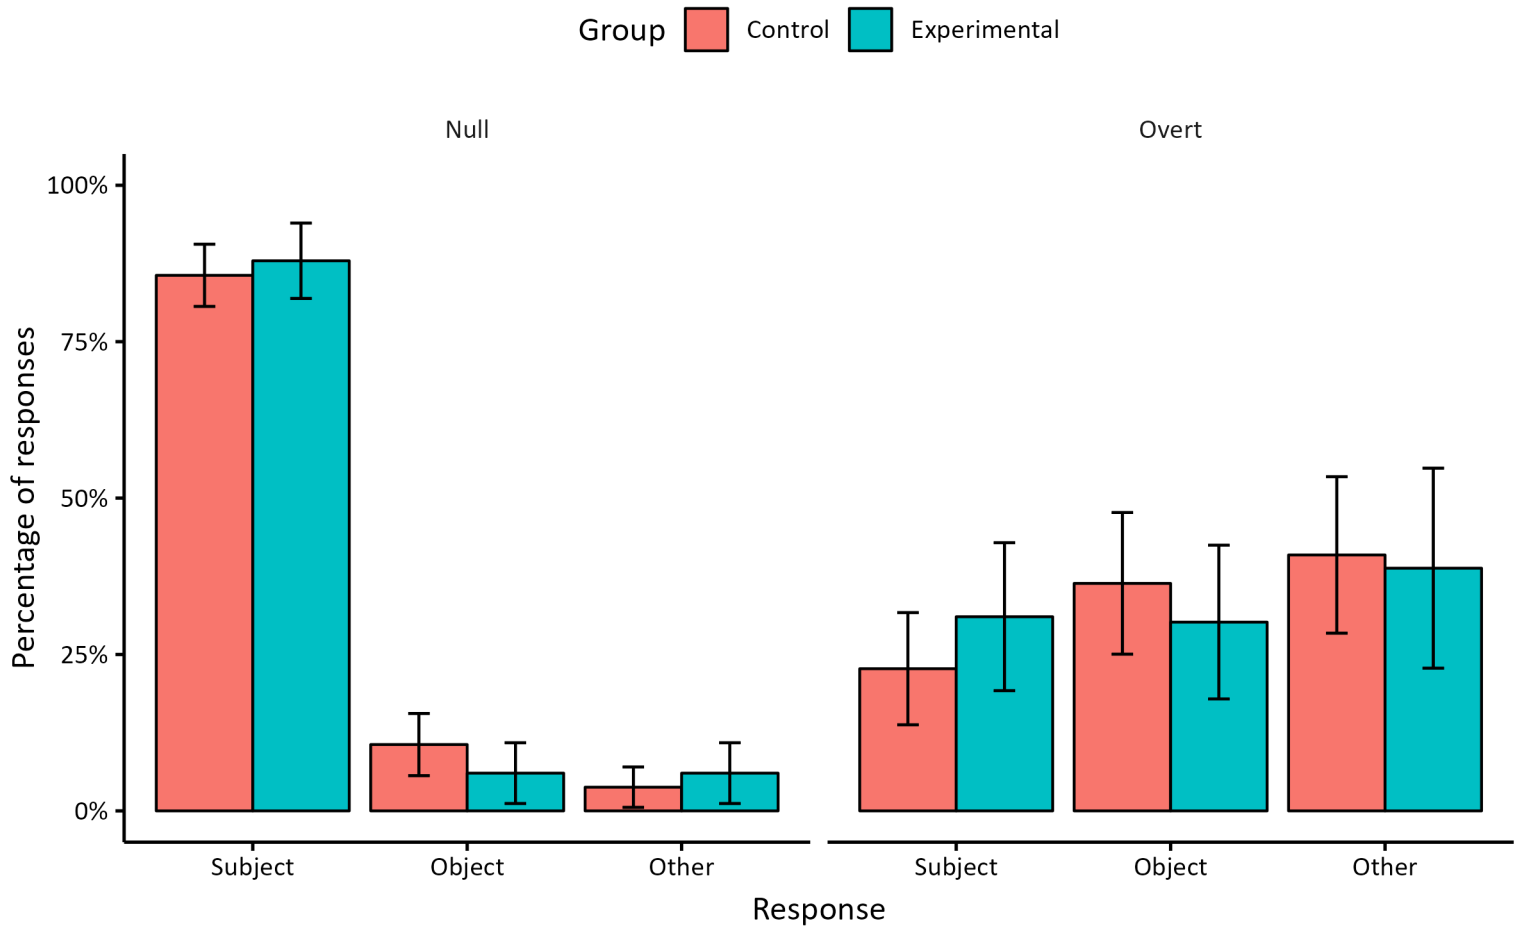

Figure D.2: Responses to globally ambiguous embedded-first pronominal items in chapter 3 with the [+TS] readings separated into 'object' and 'other'. Please note that the subject values trivially match those reported in Figure 3.1b and 'object' responses include both direct and indirect objects.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.50	0.04	169.30	$<$ **0.001**	***
Antecedent	-0.09	0.02	-3.80	$<$ **0.001**	***
Group	0.08	0.07	1.04	0.30	
Antecedent:Group	-0.10	0.05	-3.00	**0.049**	*

Table D.1: Model output for the reading times of items with overt pronouns in the fifth window for the experimental and control groups without *order*.

	Estimate	Std. Error	*t*-value	*p*-value	
Intercept	6.61	0.06	108.18	$<$ **0.001**	***
Antecedent	-0.14	0.04	-3.19	$<$ **0.01**	**
Order	-0.005	0.001	-5.72	$<$ **0.001**	***
Placement	-0.06	0.03	-2.39	**0.02**	*
Antecedent:Order	0.003	0.001	2.83	$<$ **0.01**	**

Table D.2: Model output for the reading times of items with overt pronouns in the fifth window for the experimental with *placement* as a simple effect.

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.57	0.04	153.88	$<$ **0.001**	***
Antecedent	-0.14	0.03	-4.33	$<$ **0.001**	***
Group	0.10	0.08	1.29	0.20	
Antecedent:Group	-0.19	0.06	-3.07	$<$ **0.01**	**

(a) First half of items

	Est.	Std. Error	*t*-value	*p*-value	
Intercept	6.43	0.04	165.90	$<$ **0.001**	***
Antecedent	-0.05	0.03	-1.62	0.11	
Group	0.05	0.08	0.64	0.53	
Antecedent:Group	-0.05	0.06	-0.70	0.49	

(b) Second half of items

Table D.3: Model outputs for the reading times of items with overt pronouns in the fifth window in each half of the experiment.

References

Aguilar, M., P. Ferré, J. M. Gavilán, J. A. Hinojosa & J. Demestre. 2021. The actress was on the balcony, after all: Eye-tracking locality and PR-availability effects in Spanish. *Cognition* 211. doi:10.1016/j.cognition.2021.104624.

Aguilar, M. & N. Grillo. 2021. Spanish is *not* different: On the universality of minimal structure and locality principles. *Glossa: a journal of general linguistics* 6(1). doi:10.5334/gjgl.1251.

Alonso-Ovalle, L., S. Fernández-Solera, L. Frazier & C. J. Clifton. 2002. Null vs. overt pronouns and the Topic-Focus articulation in Spanish. *Rivista di Linguistica* 14(2). 151–169. https://www.italian-journal-linguistics.com/app/uploads/2021/06/1.Alonso-Avalle.pdf.

Bates, D., M. Mächler, B. Bolker & S. Walker. 2015. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67(1). doi:10.18637/jss.v067.i01.

Baus, C., A. Costa & M. Carreiras. 2013. On the effects of second language immersion on first language production. *Acta Psychologica* 142(3). 402–409. doi:10.1016/j.actpsy.2013.01.010.

Bel, A., N. Sagarra, J. P. Comínguez & E. García-Alcaraz. 2016. Transfer and proficiency effects in l2 processing of subject anaphora. *Lingua* 184. 134–159. doi:10.1016/j.lingua.2016.07.001.

Belletti, A. 1988. The case of unaccusatives. *Linguistic inquiry* 19(1). 1–34.

Belletti, A. 2001. Inversion as focalization. In A. Hulk & J.-Y. Pollock (eds.), *Subject inversion in romance and the theory of universal grammar*, 60–90. Oxford: Oxford UniversityPress.

Belletti, A. 2004. Aspects of the low IP area. In L. Rizzi (ed.), *The structure of IP and CP. the cartography of syntactic structures*, vol. 2, New York: Oxford University Press.

Belletti, A., E. Bennati & A. Sorace. 2007. Theoretical and developmental issues in the syntax of subjects: Evidence from near-native Italian. *Natural Language & Linguistic Theory* 25(4). 657–689. doi:10.1007/s11049-007-9026-9.

Belletti, A. & C. Leonini. 2004. Subject inversion in L2 Italian. In *Eurosla yearbook*, vol. 4 95-118, John Benjamins Amsterdam/Philadelphia.

Bialystok, E., F. I. Craik & G. Luk. 2008. Lexical access in bilinguals: Effects of vocabulary size and executive control. *Journal of Neurolinguistics* 21(6). 522–538. doi:10.1016/j.jneuroling.2007.07.001.

Bianchi, V., G. Bocci & S. Cruschina. 2017. Two types of subject inversion in Italian wh-questions. *Revue Roumaine de Linguistique* 62(3). 233–252.

Bini, M. 1993. La adquisición del italiano: más allá de las propiedades sintácticas del parámetro pro-drop. In J. M. Liceras (ed.), *La lingüística y el análisis de los sistemas no nativos*, 126–139. Dovehouse.

Birdsong, D., L. Gertken & M. Amengual. 2012. Bilingual Language Profile: An Easy-to-Use Instrument to Assess Bilingualism. https://sites.la.utexas.edu/bilingual/.

Bocci, G. & L. Pozzan. 2014. Some questions (and some experimental answers) about Italian subjects: Subject positions in main and embedded questions in L1 and attrition. In C. Contemori & L. Dal Pozzo (eds.), *Inquiries in to linguistic theory and language acquisition. papers offered to Adriana Belletti.*, 28–44. CISCL Press.

Cairncross, A. & L. Dal Pozzo. 2021. Postverbal subjects in child Italian: Argument structure, discourse and the definiteness effect. *First Language* 42(3). 333–360. doi:10.1177/01427237211054247.

Cardinaletti, A. 2004. Toward a cartography of subject positions. In L. Rizzi (ed.), *The structure of CP and IP. the cartography of syntactic structures*, vol. 2, 115–165. New York: Oxford University Press.

Cardinaletti, A. & M. Starke. 1999. The typology of structural deficiency: A case study of the three classes of pronouns. In H. van Riemsdijk (ed.), *Clitics in the Languages of Europe*, vol. 5 EUROTYP, 145–233. Berlin - New York: De Gruyter Mouton.

Carminati, M. N. 2002. *The processing of Italian subject pronouns*: University of Massachusetts Amherst dissertation. https://scholarworks.umass.edu/dissertations/AAI3039345/.

Carminati, M. N. 2005. Processing reflexes of the Feature Hierarchy (Person > Number > Gender) and

implications for linguistic theory. *Lingua* 115(3). 259–285. doi:10.1016/j.lingua.2003.10.006.

Carreiras, M. & C. Clifton. 1993. Relative Clause Interpretation Preferences in Spanish and English. *Language and Speech* 36(4). 353–372. doi:10.1177/002383099303600401.

Carreiras, M. & C. Clifton. 1999. Another word on parsing relative clauses: Eyetracking evidence from Spanish and English. *Memory & Cognition* 27(5). 826–833. doi:10.3758/bf03198535.

Casalicchio, J. 2013. *Pseudorelative, Gerundi e Infiniti nelle Varietà Romanze: Affinità (solo) superficiali e corrispondenze strutturali*: Università degli Studi di Padova dissertation.

Chamorro, G. 2018. Offline interpretation of subject pronouns by native speakers of Spanish. *Glossa: a journal of general linguistics* 3(1). doi:10.5334/gjgl.256.

Chamorro, G., A. Sorace & P. Sturt. 2015a. What is the source of L1 attrition? The effect of recent L1 re-exposure on Spanish speakers under L1 attrition. *Bilingualism: Language and Cognition* 19(3). 520–532. doi:10.1017/s1366728915000152.

Chamorro, G., P. Sturt & A. Sorace. 2015b. Selectivity in l1 attrition: Differential object marking in Spanish near-native speakers of English. *Journal of Psycholinguistic Research* 45(3). 697–715. doi:10.1007/s10936-015-9372-4.

Chang, C. B. 2012. Rapid and multifaceted effects of second-language learning on first-language speech production. *Journal of Phonetics* 40(2). 249–268. doi:10.1016/j.wocn.2011.10.007.

Chomsky, N. 1995. *The Minimalist Program*. The MIT Press.

Christianson, K., A. Hollingworth, J. F. Halliwell & F. Ferreira. 2001. Thematic Roles Assigned along the Garden Path Linger. *Cognitive Psychology* 42(4). 368–407. doi:10.1006/cogp.2001.0752.

Cinque, G. 1992. The Pseudo-Relative and ACC-ing Constructions after Verbs of Perception. *University of Venice Working Papers in Linguistics* .

Contemori, C. & E. Di Domenico. 2021. Microvariation in the division of labor between null- and overt-subject pronouns: the case of Italian and Spanish. *Applied Psycholinguistics* 42(4). 997–1028. doi:10.1017/s0142716421000199.

Contemori, C. & P. E. Dussias. 2020. The processing of subject pronouns in highly proficient l2 speakers of english. *Glossa: a journal of general linguistics* 5(1). doi:10.5334/gjgl.972.

Costa, A. & M. Santesteban. 2004. Lexical access in bilingual speech production: Evidence from language switching in highly proficient bilinguals and L2 learners. *Journal of Memory and Language* 50(4). 491–511. doi:10.1016/j.jml.2004.02.002.

Cuetos, F. & D. C. Mitchell. 1988. Cross-linguistic differences in parsing: Restrictions on the use of the Late Closure strategy in Spanish. *Cognition* 30(1). 73–105. doi:10.1016/0010-0277(88)90004-2.

Cunnings, I., G. Fotiadou & I. Tsimpli. 2016. ANAPHORA RESOLUTION AND REANALYSIS DURING l2 SENTENCE PROCESSING. *Studies in Second Language Acquisition* 39(4). 621–652. doi:10.1017/s0272263116000292.

De Santo, A. & A. Y. Lee. 2022. Pseudo-relative clause effects on the online processing of Italian relative clause attachment. Presented at the 35th Conference on Human Sentence Processing. https://oxford-abstracts.s3.amazonaws.com/dde5ba24-7264-4a77-9451-a24f0eea8913.pdf.

De Vincenzi, M. 1991. *Syntactic Parsing Strategies in Italian: The Minimal Chain principle*, vol. 12 Studies in Theoretical Psycholinguistics. Dordrecht, The Netherlands: Springer Science & Business Media. doi:10.1007/978-94-011-3184-1.

De Vincenzi, M. & R. Job. 1993. Some observations on the universality of the late-closure strategy. *Journal of Psycholinguistic Research* 22. 189–206. doi:10.1007/bf01067830.

Di Domenico, E., I. Baroncini & A. Capotorti. 2020. Null and overt subject pronouns in topic continuity and topic shift: An investigation of the narrative productions of Italian Natives, Greek Natives and near-native second language speakers of Italian with Greek as a first language. *Glossa: a journal of general linguistics* 5(1). doi:10.5334/gjgl.1009.

Dussias, P. E. 2003. Syntactic Ambiguity Resolution in L2 Learners. *Studies in Second Language Acquisition* 25(4). 529–557. doi:10.1017/s0272263103000238.

Dussias, P. E. 2004. Parsing a first language like a second: The erosion of L1 parsing strategies in Spanish-English Bilinguals. *International Journal of Bilingualism* 8(3). 355–371. doi:10.1177/13670069040080031001.

Dussias, P. E. & N. Sagarra. 2007. The Effect of Exposure on Syntactic Parsing in Spanish–English Bilinguals. *Bilingualism: Language and Cognition* 10(1). 101–116. doi:10.1017/s1366728906002847.

Ecker, P. 2004. Words on the tip of the tongue: A study of lexical retrieval failures in Spanish-English bilinguals. *Southwest Journal of Linguistics* 23(2). 33–63.

Fernandes, B., A. Alexiadou, W.-Y. Chow, A. Santi & N. Grillo. 2018. Adaptation to complex cues in sentence comprehension. Presented at the 31st CUNY Conference on Human Sentence Processing.

Fernández, E. M. 2003. *Bilingual Sentence Processing*. Philadelphia, PA: John Benjamins Publishing Company. doi:10.1075/lald.29.

Filiaci, F. 2010. Null and Overt Subject Biases in Spanish and Italian: A Cross-linguistic Comparison. In C. Borgonovo, M. Español-Echevarría & P. Prévost (eds.), *Selected proceedings of the 12th Hispanic linguistics symposium*, 171–182. Somerville Massachusetts: Cascadilla Press.

Filiaci, F., A. Sorace & M. Carreiras. 2014. Anaphoric biases of null and overt subjects in Italian and Spanish: a cross-linguistic comparison. *Language, Cognition and Neuroscience* 29(7). 825–843. doi:10.1080/01690965.2013.801502.

Flege, J. E. 1987. The production of "new" and "similar" phones in a foreign language: evidence for the effect of equivalence classification. *Journal of Phonetics* 15(1). 47–65. doi:10.1016/s0095-4470(19)30537-6.

Flege, J. E. & W. Eefting. 1987. Cross-language switching in stop consonant perception and production by Dutch speakers of English. *Speech Communication* 6(3). 185–202. doi:10.1016/0167-6393(87)90025-2.

Flege, J. E. & J. Hillenbrand. 1984. Limits on phonetic accuracy in foreign language speech production. *The Journal of the Acoustical Society of America* 76(3). 708–721. doi:10.1121/1.391257.

Fodor, J. D. 1998. Learning to Parse? *Journal of Psycholinguistic Research* 27(2). 285–319. doi:10.1023/a:1023258301588.

Frazier, L. 1978. *On Comprehending Sentences: Syntactic parsing strategies*: University of Connecticut dissertation.

Frazier, L. & J. Clifton, Charles. 1995. *Construal*. MIT Press.

Fujita, H. & I. Cunnings. 2021. Reanalysis processes in non-native sentence comprehension. *Bilingualism: Language and Cognition* 24(4). 628–641. doi:10.1017/s1366728921000195.

Gade, M., M. Declerck, A. M. Philipp, A. Rey-Mermet & I. Koch. 2021. Assessing the Evidence for Asymmetrical Switch Costs and Reversed Language Dominance Effects – A Meta-Analysis. *Journal of Cognition* 4(1). doi:10.5334/joc.186.

Gargiulo, C. & J. van de Weijer. 2020. Anaphora resolution in L1 Italian in a Swedish-speaking environment before and after L1 re-immersion: A study on attrition. *Lingua* 233. 102746. doi:10.1016/j.lingua.2019.102746.

Goldrick, M. & T. H. Gollan. 2023. Inhibitory control of the dominant language: Reversed language dominance is the tip of the iceberg. *Journal of Memory and Language* 130. 104410. doi:10.1016/j.jml.2023.104410.

Gollan, T. H., R. I. Montoya, C. Cera & T. C. Sandoval. 2008. More use almost always means a smaller frequency effect: Aging, bilingualism, and the weaker links hypothesis. *Journal of Memory and Language* 58(3). 787–814. doi:10.1016/j.jml.2007.07.001.

Gollan, T. H., R. I. Montoya, C. Fennema-Notestine & S. K. Morris. 2005. Bilingualism affects picture naming but not picture classification. *Memory & Cognition* 33(7). 1220–1234. doi:10.3758/bf03193224.

van Gompel, R. P., M. J. Pickering, J. Pearson & G. Jacob. 2006. The activation of inappropriate analyses in garden-path sentences: Evidence from structural priming. *Journal of Memory and Language* 55(3).

335–362. doi:10.1016/j.jml.2006.06.004.

Greenberg, J. 1963. Some universals of grammar with particular reference to the order of meaningful elements. In J. Greenberg (ed.), *Universals of language*, 73–113. Cambridge, Massachusetts: MIT Press.

Grillo, N. 2012. Local and Universal. In V. Bianchi & C. Chesi (eds.), *Enjoy linguistics! papers offered to Luigi Rizzi on the occasion of his 60th birthday*, 234–245. Siena, Italy: CISCL Press. http://eprints.whiterose.ac.uk/110188/1/Grillo_2012_Local_and_Universal.pdf.

Grillo, N. & J. Costa. 2014. A novel argument for the Universality of Parsing principles. *Cognition* 133(1). 156–187. doi:10.1016/j.cognition.2014.05.019.

Grillo, N., J. Costa, B. Fernandes & A. Santi. 2015. Highs and Lows in English Attachment. *Cognition* 144. 116–122. doi:10.1016/j.cognition.2015.07.018.

Grillo, N. & K. Moulton. 2016. Event kinds and the pseudo relative. In . C. H. B. Prickett (ed.), *Nels 46: Proceedings of the forty-sixth annual meeting of the north east linguistics society*, vol. 2, 11–20. GLSA.

Guasti, M. T. 1993. Verb Syntax in Italian Child Grammar: Finite and Nonfinite Verbs. *Language Acquisition* 3(1). 1–40. doi:10.1207/s15327817la0301_1.

Gurel, A. 2006. L2 acquisition of pragmatic and syntactic constraints in the use of overt and null subject pronouns. In R. Slabakova, S. A. Montrul & P. Prévost (eds.), *Inquiries in linguistic development*, 259–282. John Benjamins.

Gürel, A. 2004. Selectivity in L2-induced L1 attrition: a psycholinguistic account. *Journal of Neurolinguistics* 17(1). 53–78. doi:10.1016/s0911-6044(03)00054-x.

Gürel, A. & G. Yılmaz. 2011. Restructuring in the L1 Turkish grammar. *Language, Interaction and Acquisition* 2(2). 221–250. doi:10.1075/lia.2.2.03gur.

Hacohen, A. & J. Schaeffer. 2007. Subject realization in early Hebrew/English bilingual acquisition: The role of crosslinguistic influence. *Bilingualism: Language and Cognition* 10(3). 333–344. doi:10.1017/s1366728907003100.

Harley, H. & E. Ritter. 2002. Person and Number in Pronouns: A Feature-Geometric Analysis. *Language* 78(3). 482–526. doi:10.1353/lan.2002.0158.

Hemforth, B., S. Fernandez, C. Clifton, L. Frazier, L. Konieczny & M. Walter. 2015. Relative clause attachment in German, English, Spanish and French: Effects of position and length. *Lingua* 166. 43–64. doi:10.1016/j.lingua.2015.08.010.

Hendriks, P. 2014. *Asymmetries between Language Production and Comprehension*. Springer Netherlands. doi:10.1007/978-94-007-6901-4.

Hendriks, P., C. Koster & J. C. Hoeks. 2014. Referential choice across the lifespan: why children and elderly adults produce ambiguous pronouns. *Language, Cognition and Neuroscience* 29(4). 391–407. doi:10.1080/01690965.2013.766356.

Hilles, S. 1986. Interlanguage and the pro-drop parameter. *Interlanguage studies bulletin (Utrecht)* .

Holbrook, A. 2008. Acquiescence response bias. In P. Lavrakas (ed.), *Encyclopedia of survey research methods*, Sage Publications, Inc. doi:10.4135/9781412963947.

Hyams, N. 1986. *Language acquisition and the theory of parameters*. Dordrecht: D. Reidel.

Ivanova, I. & A. Costa. 2008. Does bilingualism hamper lexical access in speech production? *Acta Psychologica* 127(2). 277–288. doi:10.1016/j.actpsy.2007.06.003.

Jacob, G. & C. Felser. 2016. Reanalysis and semantic persistence in native and non-native garden-path recovery. *Quarterly Journal of Experimental Psychology* 69(5). 907–925. doi:10.1080/17470218.2014.984231.

Jegerski, J., B. VanPatten & G. D. Keating. 2011. Cross-linguistic variation and the acquisition of pronominal reference in L2 Spanish. *Second Language Research* 27(4). 481–507. doi:10.1177/0267658311406033.

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

Kaltsa, M., I. M. Tsimpli & J. Rothman. 2015. Exploring the source of differences and similarities in L1

attrition and heritage speaker competence: Evidence from pronominal resolution. *Lingua* 164. 266–288. doi:10.1016/j.lingua.2015.06.002.

Keating, G. D., J. Jegerski & B. Vanpatten. 2016. Online processing of subject pronouns in monolingual and heritage bilingual speakers of Mexican Spanish. *Bilingualism: Language and Cognition* 19(1). 36–49. doi:10.1017/s1366728914000418.

Keating, G. D., B. VanPatten & J. Jegerski. 2011. WHO WAS WALKING ON THE BEACH? *Studies in Second Language Acquisition* 33(2). 193–221. doi:10.1017/s0272263110000732.

Kim, J., S. Montrul & J. Yoon. 2009a. Dominant language influence in acquisition and attrition of binding: Interpretation of the Korean reflexive *caki*. *Bilingualism: Language and Cognition* 13(1). 73–84. doi:10.1017/s136672890999037x.

Kim, J.-H., S. Montrul & J. Yoon. 2009b. Binding Interpretations of Anaphors by Korean Heritage Speakers. *Language Acquisition* 16(1). 3–35. doi:10.1080/10489220802575293.

Kimball, J. 1973. Seven principles of surface structure parsing in natural language. *Cognition* 2(1). 15–47. doi:10.1016/0010-0277(72)90028-5.

Köpke, B. & D. Genevska-Hanke. 2018. First Language Attrition and Dominance: Same Same or Different? *Frontiers in Psychology* 9. doi:10.3389/fpsyg.2018.01963.

Lee, A. Y. & A. De Santo. 2022. Online processing effects of Pseudo Relative Clause Availability in Italian Attachment Ambiguity. Presented at the 46th Penn Linguistics Conference. https://www.ling.upenn.edu/Events/PLC/plc46/PLC46Abstracts.pdf.

Lenth, R. V. 2022. *emmeans: Estimated Marginal Means, aka Least-Squares Means*. https://CRAN.R-project.org/package=emmeans. R package version 1.7.5.

Liceras, J. & L. Díaz. 1999. Topic-drop versus pro-drop: null subjects and pronominal subjects in the Spanish L2 of Chinese, English, French, German and Japanese speakers. *Second Language Research* 15(1). 1–40.

Liceras, J. M. 1989. On some properties of the "pro-drop" parameter: looking for missing subjects in non-native Spanish. In S. M. Gass & J. Schachter (eds.), *Linguistic perspectives on second language*, Cambridge University Press.

Linck, J. A., J. F. Kroll & G. Sunderman. 2009. Losing Access to the Native Language While Immersed in a Second Language: Evidence for the Role of Inhibition in Second-Language Learning. *Psychological Science* 20(12). 1507–1515. doi:10.1111/j.1467-9280.2009.02480.x.

Lorusso, P., C. Caprin & M. T. Guasti. 2004. Overt Subject Distribution in Early Italian Children. In A. Brugos, M. R. Clark-Cotton & S. Ha (eds.), *BUCLD 29 online proceedings supplement*, https://www.bu.edu/bucld/files/2011/05/29-LorussoBUCLD2004.pdf.

Lozano, C. 2006. The development of the syntax–discourse interface: Greek learners of Spanish. In . L. E. V. Torrens (ed.), *The acquisition of syntax in romance languages*, 371–399. Amsterdam: John Benjamins.

Lozano, C. 2017. The Development of Anaphora Resolution at the Syntax-Discourse Interface: Pronominal Subjects in Greek Learners of Spanish. *Journal of Psycholinguistic Research* 47(2). 411–430. doi:10.1007/s10936-017-9541-8.

Major, R. C. 1992. Losing English as a First Language. *The Modern Language Journal* 76(2). 190–208. doi:10.1111/j.1540-4781.1992.tb01100.x.

Malyutina, S. & D.-B. den Ouden. 2016. What is it that lingers? Garden-path (mis)interpretations in younger and older adults. *Quarterly Journal of Experimental Psychology* 69(5). 880–906. doi:10.1080/17470218.2015.1045530.

Margaza, P. & A. Bel. 2006. Null subjects at the syntax–pragmatics interface: Evidence from Spanish interlanguage of Greek speakers. In M. G. O'Brien, C. Shea & J. Archibald (eds.), *Proceedings of GASLA 2006*, 88–97. Somerville Massachusetts: Cascadilla Press.

Martín-Villena, F. 2023. *L1 MORPHOSYNTACTIC ATTRITION AT THE EARLY STAGES: EVIDENCE FROM PRODUCTION, INTERPRETATION, AND PROCESSING OF SUBJECT REFERRING EXPRES-*

SIONS IN L1 SPANISH-L2 ENGLISH INSTRUCTED AND IMMERSED BILINGUALS: Universidad de Grenada dissertation.

Matuschek, H., R. Kliegl, S. Vasishth, H. Baayen & D. Bates. 2017. Balancing type i error and power in linear mixed models. *Journal of Memory and Language* 94. 305–315. doi:10.1016/j.jml.2017.01.001.

Mayr, R., S. Price & I. Mennen. 2012. First language attrition in the speech of Dutch–English bilinguals: The case of monozygotic twin sisters. *Bilingualism: Language and Cognition* 15(4). 687–700. doi:10.1017/s136672891100071x.

Meuter, R. F. I. & A. Allport. 1999. Bilingual Language Switching in Naming: Asymmetrical Costs of Language Selection. *Journal of Memory and Language* 40(1). 25–40. doi:10.1006/jmla.1998.2602.

Mitchell, D. C. & F. Cuetos. 1991. The origins of parsing strategies. *Current issues in natural language processing* 1–12.

Montrul, S. 2004. Subject and object expression in Spanish heritage speakers: A case of morphosyntactic convergence. *Bilingualism: Language and Cognition* 7(2). 125–142. doi:10.1017/s1366728904001464.

Montrul, S., R. Bhatt & R. Girju. 2015. Differential object marking in Spanish, Hindi, and Romanian as heritage languages. *Language* 91(3). 564–610. doi:10.1353/lan.2015.0035.

Montrul, S. & C. R. Louro. 2006. Beyond the syntax of the Null Subject Parameter: a look at the discourse-pragmatic distributionof null and overt subjects by L2 learners of Spanish. In L. Escobar & V. Torrens (eds.), *The acquisition of syntax in romance languages*, John Benjamins.

Moulton, K. & N. Grillo. 2015. Pseudo Relatives: Big and Direct. In B. . Ozyildiz (ed.), *Proceedings of 45 north eastern linguistic society*, 193–202. Graduate Student Linguistics Association.

Otheguy, R., A. C. Zentella & D. Livert. 2007. LANGUAGE AND DIALECT CONTACT IN SPANISH IN NEW YORK: TOWARD THE FORMATION OF A SPEECH COMMUNITY. *Language* 83(4). 770–802.

Papadopoulou, D. & H. Clahsen. 2003. PARSING STRATEGIES IN L1 AND L2 SENTENCE PROCESSING. *Studies in Second Language Acquisition* 25(4). 501–528. doi:10.1017/s0272263103000214.

Paradis, J. & S. Navarro. 2003. Subject realization and crosslinguistic interference in the bilingual acquisition of Spanish and English: what is the role of the input? *Journal of Child Language* 30(2). 371–393. doi:10.1017/s0305000903005609.

Paradis, M. 1993. Linguistic, psycholinguistic, and neurolinguistic aspects of "interference" in bilingual speakers: The activation threshold hypothesis. *International Journal of Psycholinguistics* 9(2). 133–145.

Paradis, M. 2007. L1 attrition features predicted by a neurolinguistic theory of bilingualism. In B. Köpke, M. S. Schmid, M. Keijzer & S. Dostert (eds.), *Language attrition*, 121–134. The Netherlands: John Benjamins Publishing Company.

Pesetsky, D. 1987. Wh-in-Situ: movement and unselective binding. In t. A. Reuland E. (ed.), *The representation of (in)definiteness*, 98–129. MIT Press.

Pierce, A. E. 1992. *Language Acquisition and Syntactic Theory: A Comparative Analysis of French and English Child Grammars* Studies in Theoretical Psycholinguistics. Springer Netherlands.

Pinto, M. 2006. Subject pronouns in bilinguals. In V. Torrens & L. Escobar (eds.), *The acquisition of syntax in romance languages*, 331–350. John Benjamins.

Pires, A. & J. Rothman. 2009. Disentangling sources of incomplete acquisition: An explanation for competence divergence across heritage grammars. *International Journal of Bilingualism* 13(2). 211–238. doi:10.1177/1367006909339806.

Platzack, C. 1987. The Scandinavian Languages and the Null-Subject Parameter. *Natural Language & Linguistic Theory* 5(3). 377–401.

Portocarrero, J., R. Burright & P. Donovick. 2007. Vocabulary and verbal fluency of bilingual and monolingual college students. *Archives of Clinical Neuropsychology* 22(3). 415–422. doi:10.1016/j.acn.2007.01.015.

Pozniak, C., B. Hemforth, Y. Haendler, A. Santi & N. Grillo. 2019. Seeing events vs. entities: The processing advantage of Pseudo Relatives over Relative Clauses. *Journal of Memory and Language* 107. 128–151.

doi:10.1016/j.jml.2019.04.001.

Pozzan, L. & J. C. Trueswell. 2015. Second language processing and revision of garden-path sentences: a visual word study. *Bilingualism: Language and Cognition* 19(3). 636–643. doi:10.1017/s1366728915000838.

Pérez-Leroux, A. T. & W. R. Glass. 1999. Null anaphora in Spanish second language acquisition: probabilistic versus generative approaches. *Second Language Research* 15(2). 220–249.

R Core Team. 2022. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing Vienna, Austria. https://www.R-project.org/.

Radford, A. 1975. Pseudo-Relatives and the Unity of Subject-Raising. *Archivum Linguisticum* 6. 32–64.

Reinhart, T. 1981. Pragmatics and linguistics: An analysis of sentence topics. *Philosophica* 1. 53–94.

Rizzi, L. 1982. *Issues in Italian Syntax*. Dordrecht, Netherlands: Foris Publications.

Rizzi, L. 1986. Null Objects in Italian and the Theory of *pro*. *Linguistic Inquiry* 17(3). 501–557. http://www.jstor.org/stable/4178501.

Rizzi, L. 1993. Some Notes on Linguistic Theory and Language Development: The Case of Root Infinitives. *Language Acquisition* 3(4). 371–393. doi:10.1207/s15327817la0304_2.

Rizzi, L. 2001. On the position Int(errogative) in the left periphery of the clause. In G. Cinque & G. Salvi (eds.), *Current studies in Italian syntax. essays offered to Lorenzo Renzi*, 287–296. Brill.

Rizzi, L. 2006. Selective Residual V-2 in Italian Interrogatives. In P. Brandt & E. Fuß (eds.), *Form, structure and grammar: A festschrift presented to günther grewendorf on occasion of his 60th birthday.*, 229–242. De Gruyter, Inc.

Roberts, L., M. Gullberg & P. Indefrey. 2008. ONLINE PRONOUN RESOLUTION IN L2 DISCOURSE: L1 Influence and General Learner Effects. *Studies in Second Language Acquisition* 30(3). 333–357. doi:10.1017/s0272263108080480.

Rothman, J. 2007. Heritage speaker competence differences, language change, and input type: Inflected infinitives in Heritage Brazilian Portuguese. *International Journal of Bilingualism* 11(4). 359–389. doi:10.1177/13670069070110040201.

Rothman, J. 2009. Pragmatic deficits with syntactic consequences?: L2 pronominal subjects and the syntax–pragmatics interface. *Journal of Pragmatics* 41(5). 951–973. doi:10.1016/j.pragma.2008.07.007.

Schmid, M. S. & E. Dusseldorp. 2010. Quantitative analysesin a multivariate studyof language attrition: the impactof extralinguistic factors. *Second Language Research* 21(1). 125–160.

Schmid, M. S. & S. Jarvis. 2014. Lexical access and lexical diversity in first language attrition. *Bilingualism: Language and Cognition* 17(4). 729–748. doi:10.1017/s1366728913000771.

Schmid, M. S. & B. Köpke (eds.). 2019. *The Oxford Handbook of Language Attrition*. Oxford University Press. doi:10.1093/oxfordhb/9780198793595.001.0001.

Schmid, M. S. & G. Yilmaz. 2021. Lexical Access in L1 Attrition–Competition versus Frequency: A Comparison of Turkish and Moroccan Attriters in the Netherlands. *Applied Linguistics* 42(5). 878–904. doi:10.1093/applin/amab006.

Şener, S. & D. Takahashi. 2010. ELLIPSIS OF ARGUMENTS IN JAPANESE AND TURKISH. *Nanzan Linguistics* 6. 79–99.

Serratrice, L. 2007. Cross-linguistic influence in the interpretation of anaphoric and cataphoric pronouns in English–Italian bilingual children. *Bilingualism: Language and Cognition* 10(3). 225–238. doi:10.1017/s1366728907003045.

Serratrice, L., A. Sorace, F. Filiaci & M. Baldo. 2012. Pronominal objects in English–Italian and Spanish–Italian bilingual children. *Applied Psycholinguistics* 33(4). 725–751. doi:10.1017/s0142716411000543.

Serratrice, L., A. Sorace & S. Paoli. 2004. Crosslinguistic influence at the syntax–pragmatics interface: Subjects and objects in English–Italian bilingual and monolingual acquisition. *Bilingualism: Language and Cognition* 7(3). 183–205. doi:10.1017/s1366728904001610.

Shen, X. 2006. *Late assignment of syntax theory: Evidence from Chinese and English*: University of Exeter

dissertation.

Shin, N. L. & H. S. Cairns. 2012. The Development of NP Selection in School-Age Children: Reference and Spanish Subject Pronouns. *Language Acquisition* 19(1). 3–38. doi:10.1080/10489223.2012.633846.

Silva-Corvalán, C. 2014. *Bilingual language acquisition: Spanish and English in the first six years*. Cambridge University Press. doi:10.1017/cbo9781139162531.

Silverstein, M. 1985. Hierarchy of features and ergativity. In P. Muysken & H. van Riemsdijk (eds.), *Features and projections*, 163–232. Foris Publications.

Slattery, T. J., P. Sturt, K. Christianson, M. Yoshida & F. Ferreira. 2013. Lingering misinterpretations of garden path sentences arise from competing syntactic representations. *Journal of Memory and Language* 69(2). 104–120. doi:10.1016/j.jml.2013.04.001.

Slioussar, N., D. Antropova & D. Chernova. 2022. To forget or not to forget: processing gender, number and case features. Presented at AMLaP 2022.

Sorace, A. 2000. Differential Effects of Attrition in the L1 Syntax of Near-native L2 Speakers. In S. C. Howell, S. A. Fish & T. Keith-Lucas (eds.), *Proceedings of the 24th annual Boston University conference on language development*, vol. 2, 719–725. Somerville, Massachusetts: Cascadilla Press.

Sorace, A. 2011. Pinning down the concept of "interface" in bilingualism. *Linguistic Approaches to Bilingualism* 1(1). 1–33. doi:10.1075/lab.1.1.01sor.

Sorace, A. 2016. Referring expressions and executive functions in bilingualism. *Linguistic Approaches to Bilingualism* 6(5). 669–684. doi:10.1075/lab.15055.sor.

Sorace, A. & F. Filiaci. 2006. Anaphora resolution in near-native speakers of Italian. *Second Language Research* 22(3). 339–368. doi:10.1191/0267658306sr271oa.

Sorace, A., L. Serratrice, F. Filiaci & M. Baldo. 2009. Discourse conditions on subject pronoun realization: Testing the linguistic intuitions of older bilingual children. *Lingua* 119(3). 460–477. doi:10.1016/j.lingua.2008.09.008.

Sturt, P. 2007. Semantic re-interpretation and garden path recovery. *Cognition* 105(2). 477–488. doi:10.1016/j.cognition.2006.10.009.

Tabor, W., B. Galantucci & D. Richardson. 2004. Effects of merely local syntactic coherence on sentence processing. *Journal of Memory and Language* 50(4). 355–370. doi:10.1016/j.jml.2004.01.001.

Torregrossa, J., C. Bongartz & I. M. Tsimpli. 2019. Bilingual reference production. *Linguistic Approaches to Bilingualism* 9(4-5). 569–599. doi:10.1075/lab.17026.tor.

Tsimpli, I. & A. Roussou. 1991. Parameter-Resetting in L2? *UCL Working Papers in Linguistics* .

Tsimpli, I. & A. Sorace. 2006. Differentiating Interfaces: L2 Performance in Syntax-Semantics and Syntax-Discourse Phenomena. In D. Bamman, T. Magnitskaia & C. Zaller (eds.), *Proceedings of the 3th annual Boston University Conference on Language Development*, vol. 2, 653–664. Somerville, MA: Cascadilla Press.

Tsimpli, I., A. Sorace, C. Heycock & F. Filiaci. 2004. First language attrition and syntactic subjects: A study of Greek and Italian near-native speakers of English. *International Journal of Bilingualism* 8(3). 257–277. doi:10.1177/13670069040080030601.

Valian, V. 1991. Syntactic subjects in the early speech of American and Italian children. *Cognition* 40(1-2). 21–81. doi:10.1016/0010-0277(91)90046-7.

Valian, V. & Z. Eisenberg. 1996. The development of syntactic subjects in Portuguese-speaking children. *Journal of Child Language* 23(1). 103–128. doi:10.1017/s0305000900010114.

Vogelzang, M., F. Foppolo, M. T. Guasti, H. van Rijn & P. Hendriks. 2019. Reasoning about alternative forms is costly: The processing of null and overt pronouns in Italian using pupillary responses. *Discourse Processes* 57(2). 158–183. doi:10.1080/0163853x.2019.1591127.

Vogelzang, M., M. T. Guasti, H. Rijn & P. Hendriks. 2021. How Children Process Reduced Forms: A Computational Cognitive Modeling Approach to Pronoun Processing in Discourse. *Cognitive Science*

45(4). doi:10.1111/cogs.12951.

Wexler, K. 1994. Optional infinitives, head movement and the economy of derivations. In D. Lightfoot & N. Hornstein (eds.), *Verb movement*, 305–350. Cambridge University Press.

White, L. 1985. THE "PRO-DROP" PARAMETER IN ADULT SECOND LANGUAGE ACQUISITION. *Language Learning* 35(1). 47–61. doi:10.1111/j.1467-1770.1985.tb01014.x.

White, L. 2011. Second language acquisition at the interfaces. *Lingua* 121(4). 577–590. doi:10.1016/j.lingua.2010.05.005.

Yilmaz, G. & M. S. Schmid. 2012. L1 accessibility among Turkish-Dutch bilinguals. *The Mental Lexicon* 7(3). 249–274. doi:10.1075/ml.7.3.01yil.

Zagar, D., J. Pynte & S. Rativeau. 1997. Evidence for Early Closure Attachment on First Pass Reading Times in French. *The Quarterly Journal of Experimental Psychology Section A* 50(2). 421–438. doi:10.1080/713755715.

Zehr, J. & F. Schwarz. 2018. *PennController for Internet Based Experiments (IBEX)*. doi:10.17605/OSF.IO/MD832.