Learnability in Romance: How Indirect Input Helps Children Acquire The contrast between Null and Overt Subjects

Hannah Forsythe*, Daniel Greeson**, and Cristina Schmitt**

1. Introduction

Since the seminal work by Hyams (1986/2012), the acquisition of the so-called null subject parameter has received a fair amount of attention. The field has learned a lot by comparing children's null subject use across languages that do and do not allow null subjects. For instance, while children initially produce null subjects independent of the target language, the rate of null subject use is much lower among children exposed to non pro-drop languages (Valian 1990, inter alia). We have also learned that not all pro-drop languages are created equal: in Italian-style languages, verbal agreement either licenses *pro* or is itself the pronominal element that serves as the subject; in topic-drop languages like Mandarin (and arguably Brazilian Portuguese and Finnish; see Barbosa 2014), null subjects (and objects) are licensed by an operator linked to a topic position.

In most of these languages, children seem to converge quickly on the target (Yang 2002 and references therein). Nevertheless, it would be a mistake to assume that acquisition ends as soon as the target parameter is set. Children must still learn the conditions governing the appropriate use of null and overt subjects, and this is a complex process because the distribution of null and overt subjects is not categorical. Rather, it depends on subtle syntactic, semantic and discourse factors which vary across languages and language varieties.

In this paper, we ask how children learn the conditions governing the use of null and overt subject pronouns in Mexican Spanish and use it in comprehension. We argue that it is more efficient for children to acquire this contrast by initially observing the realization of 1st and 2nd person subjects in their input, and we provide evidence that by 4½ years of age, children both produce this contrast and generalize it to the comprehension of ambiguous 3rd person pronominal subjects.

2. Linguistic background (adult grammar)

For the purposes of this paper we assume that both null and overt pronominal subjects are variables, interpreted by an assignment function (Schlenker 1999, Sauerland 2008, Charnavel 2017), and that a pronoun's phi-features introduce presuppositions restricting the value of this variable. This is illustrated in (1) from Büring (2011), where the assignment function g is essentially an ordered list of referents that takes the index of the pronoun, i, and returns the i-th member of that list, and the feminine and singular features of the pronoun cause the output of g to be undefined unless their presuppositions are satisfied.

(1) $[\![she_i]\!]^g = g(i)$
    where $g(i)$ returns the i-th member of g,
    and $g(i)$ is a singular female, else undefined.

On this view, the task of the listener is to construct the assignment function based on what he or she believes is the intended referent. This is a very subtle process that depends on a multitude of considerations, including world knowledge and coherence relations between the events in the discourse (Kehler 2002). For example, in (2a) the hearer is unlikely to pick out *Bill* as the referent of *he* despite both Bill and John being syntactically viable options, due to real-world knowledge about the nature of apologies. On the other hand, if the coherence relation between the two events changes, as in (2b), the preference is reversed.

(2) a. John apologized to Bill because he had been rude.
    b. John apologized to Bill in spite of the fact that he had been rude.

Speakers of so-called consistent null subject languages (Italian, Spanish, etc.) have one extra property to consider in determining the intended referent of a pronominal subject. They must consider the division of labor between the null and the overt pronoun in structures where both are grammatical. Many proposals have been made to explain this distinction (Carminati 2002, Blackwell & Quesada 2012, and many others) and in all cases, they boil down to some version of prominence: the more prominent the antecedent, the more likely the subject is to be null. Proposals diverge on the exact metric that should be used to determine the hierarchy of potential antecedents (see de la Fuente 2015 for a review). There is, however, a consensus that the null subject is more likely to have an antecedent in subject position, relative to an overt pronominal subject, all things being equal (though of course in many cases all things are not equal). In (3) for example, the null is more likely to refer to *Juan* than is the overt pronoun *él* (he).

(3) a. *Juan llamó a Pedro cuando Ø estaba en casa.*
       Juan called Pedro when ***pro*** was at home.
    b. *Juan llamó a Pedro cuando él estaba en casa.*
       Juan called Pedro when **he** was   at home.

The subject-antecedent reading is commonly referred to as the "same-reference" reading because the same referent, *Juan*, is referred to in subject

position both times; all other readings tend to be grouped under the term "switch-reference." The expectation that null subjects maintain reference while overt pronominal subjects shift reference is revealed not only by speakers' antecedent preferences (Carminati 2002, Alonso-Ovalle *et al*. 2002), but also by findings from adult processing, which provide convincing evidence of reduced acceptability and longer reading times when these expectations are violated (Filiaci *et al*. 2013, Jegerski *et al*. 2011, Keating *et al*. 2016).

It is important to point out that antecedent position does not uniquely determine the distribution of overt and null pronominal subjects. Both (3a) and (3b) can grammatically be used to refer to the preceding subject, the object, or an extra-sentential referent, and what's more, any of these readings can be made more or less likely by the surrounding discourse. For example, were it made known to the listener that Juan was calling Pedro *in order to find out where Pedro was at that moment*, then the object reading would be favored over the subject reading in both (3a) and (3b), although the object bias would be slightly weaker for (3a). Moreover, speakers of different consistent pro-drop languages show some variation. Processing studies show that some varieties more strongly associate the overt variant with a non-subject antecedent compared to others (Filiaci *et al*. 2013, de la Fuente 2015), while sociolinguistic studies show that linguistic and social factors can condition overt subject use differently across language varieties (Otheguy *et al*. 2007; Carvalho *et al.* 2015).

How do children acquire this probabilistic association between pronoun realization and pronoun reference? Even assuming that the tendency to treat subject antecedents as the most prominent element, and therefore to associate them with more reduced pronominal forms, is in some sense universal, there is still learning to be done. Children must learn how this association is implemented in their particular language variety, and they must learn to properly restrict the size of the window from which to draw potential antecedents. This is something far from trivial, as illustrated by example (4) below, where there are in fact two subject antecedents (*the dress, my purse*), and the listener (in this case Speaker A) must decide whether to include one or both as potential antecedents for the pronoun *it*.

(4) Speaker A: I've decided which one to buy. After all, this dress is on sale and my purse is the same color.
    Speaker B: Are you sure it's the right color?

2. The learning problem

Before we propose a hypothesis for how this contrast may be acquired, it is important to make the learning problem clear. Spanish licenses a variety of entity-denoting constituents in subject position, including definite and, indefinite noun phrases, overt pronouns, and null pronouns. The learner's task is to decide when to use each one. In the 1st and 2nd persons, the problem is limited to the alternation between null and overt pronouns, since other DP subjects are rare (limited to 1st person plural cases like *las mamás somos inteligentes* "we mothers are-1P smart"). If children can use this alternation to identify null and overt pronouns as

members of the same set of alternatives, the next important step is to identify which one is the unmarked form—that is, the one that surfaces when conditions for use of the more marked form fail to be satisfied. Given that in pro-drop languages null subjects overwhelmingly outnumber overt pronominal subjects, sheer frequency could help the learner decide on the null subject as the default. (Table 1 illustrates the stark contrast in frequency between null and overt pronominal subjects in a number of Spanish varieties.) The crux of the learning problem therefore lies in identifying which environments trigger the use of the overt pronoun.

What could help the learner identify switch-reference as a crucial environmental trigger for the overt pronoun? This is not an easy problem to solve because identifying same- and switch-reference environments requires the learner to know in advance what each pronoun refers to—something that depends on a multitude of syntactic and pragmatic factors, as discussed above. In other words, in order to identify the constraints on null and overt pronoun use, children might first need to acquire the constraints on pronoun resolution itself.

Table 1: Proportion of null and overt subjects across varieties of Spanish.

	% overt pronouns
Barranquilla (Orozco & Guy 2008)	35.7%
Madrid (Enríquez 1984)	21%
Mexico City(Lastra & Butragueño 2015)	21.7%
NYC 1st Generation(Shin & Erker 2015)	21%
Yucatán (Michnowicz 2015)	20%

Charnavel (2017) argues that 1st and 2nd person pronouns have the same formal representation as 3rd person pronouns—all three depend on an assignment function linking their index to the intended referent. The reason for the interpretive differences that researchers have observed between 1st and 2nd person on the one hand and 3rd person on the other, is due to the fact that, for the former, it is simply a lot easier to infer the speaker's intended referent. For first person, the intended referent is nearly always the speaker, and for second person it is almost always the addressee. This is why, for instance, 1st and 2nd person often resist the kinds of bound readings available for 3rd person pronouns, even though it has been known since Kratzer (2009) that bound readings are possible in certain environments.

Charnavel's analysis suggests two things: (i) that it should be much easier for children to identify examples of 1st and 2nd person pronouns appearing in same-reference and switch-reference contexts, and (ii) that there is nothing preventing children from automatically generalizing knowledge they have about 1st and 2nd person subject pronouns to the 3rd person. This leads us to hypothesize that children acquire the null/overt subject contrast by first observing the realization

of 1st and 2nd person subjects in their input. Once the association between overt realization and switch-reference is formed, it should automatically generalize to their production and comprehension of 3rd person pronominal subjects. This learning path is summarized below.

(5) Proposed learning path:
    a. Step 1—Association: Track the realization of 1st and 2nd person subjects in same- and switch-reference contexts and associate switch-reference readings with an increased probability of overt realization.
    b. Step 2—Generalization: Generalize this association to the production/comprehension of *all* subject pronouns—1st, 2nd and 3rd person.

In other work (Forsythe *et al.*, submitted), we provide production data supporting this claim for Mexico City Spanish. Section 4 briefly presents this data and shows how it is consistent with the proposal in (5). What this production data cannot show, however, is that the generalization is complete and automatic. That is, as soon as children show signs of having acquired the contrast between null and overt 1st and 2nd person subjects, this knowledge should be available to them when interpreting 3rd person pronouns. Section 5 therefore presents a pronoun resolution task probing Mexican children's ability to use the null/overt contrast to guide their interpretation of grammatically ambiguous 3rd person subject pronouns.

Before presenting out production and comprehension results from Mexico City Spanish, we first review what is already known about children's acquisition of this contrast in pro-drop languages more generally.

3. Acquisition background

What we know so far about the developmental path comes from studies of spontaneous production, felicity judgments, and pronoun resolution tasks in a variety of languages. We concentrate here on pro-drop languages with rich agreement like Spanish, Catalan, Italian, and Greek. Children learning these languages must associate switch-reference contexts with an increased rate of overt pronoun realization, relative to same-reference contexts, although the strength of this association may vary across languages (Filiaci *et al.* 2013, de la Fuente 2015).

Production studies from the acquisition literature find that children acquiring Spanish and Catalan begin to produce overt subjects before age 2 (Grinstead 2004, Bel 2003). However, overt pronominal subjects remain infrequent. Sociolinguistic studies that include child subjects corroborate the low frequency of overt pronominal subjects well into middle and late childhood. In Mexican Spanish, first graders (ages 6-7) overtly realize pronouns at a rate of 6-8% (Shin 2012, 2016), and this rate gradually increases to about 10% in 5th grade (ages 10-11)—far less than the 18-22% rate found among Mexican adults (Lastra & Butragueño 2015, Shin & Otheguy 2013, Shin & Erker 2015). Despite being infrequent, however, overt pronouns are not randomly distributed. Shin (2016) finds that overt realization is positively associated with switch-reference among

even the youngest age group (6-7 years).

Felicity judgment tasks show that children are much more accepting than adults of null subjects appearing in switch-reference contexts as well as of overt subjects appearing in same-reference contexts. Sorace *et al.* (2009) asked bilingual children, monolingual children, and adult speakers of Italian to judge which of two speakers produced a better description of an event, with one speaker producing a null subject (6a) and the other an overt subject (6b). Crucially, the event was manipulated such that either *Minnie* herself had fallen (a same-reference context) or another person had fallen (a switch-reference context). Monolingual Italian-acquiring children ages 6-7 and 8-10 were just as likely as adults to prefer the overt pronoun in the switch-reference condition; however, the 6–7-year-olds were not as likely as adults to prefer the null pronoun in the same-reference condition. Bilinguals were overall less adult-like than their monolingual peers—even those whose other language was Spanish, another pro-drop language. Shin & Cairns (2012) used a similar methodology to probe the preferences of Mexican Spanish-acquiring children from ages 6 to 15. They found a qualitatively similar but quantitatively slower developmental trajectory. In the switch-reference condition, 8-10-year-olds showed a significant preference for an overt subject, although the strength of this preference did not match adults until age 14–15. In the same-reference condition, even the oldest children failed to show a significant preference for the null subject. Finally, studies among L2 learners have shown an even more pronounced version of this pattern of non-target behavior (Belletti *et al.* 2007, Montrul 2004; White 2011, a.o.).

(6) a. *Minnie ha detto che ø   è caduta.*
        Minnie has said that (**she**) has fallen.
      b. *Minnie ha detto che lei è caduta.*
         Minnie has said that **she** has fallen.

These results suggest a very gradual developmental path; however, they may underestimate what children actually know about the null/overt contrast. The tasks employed in these studies require the listener to hold two utterances in short-term memory while making a metalinguistic judgment about them, something rather taxing for children. Papadopoulou *et al.* (2015) used a less taxing method to probe Greek-acquiring children's sensitivity to the null/overt distinction from ages 6 to 11 and found more adult-like performance. Participants listened to a sentence like (7a) while simultaneously viewing a picture that corresponded to either a same-reference interpretation (*pro* = the old man), or one of two switch-reference interpretations (object interpretation: *pro* = his grandchild; other interpretation: *pro* = another person). At the end of each sentence, the task was simply to judge whether the sentence matched the picture. Additionally, the audio was self-paced: participants would press the spacebar to hear each phrase of the sentence, and their listening times were measured. A second experiment used the same methodology with overt pronouns, as in (7b).

(7) a. *O papús   millúse   ðinatá ston egonó tu   ótan ø ðjávaze ena vivlío.*
        The old-man spoke-3Sg loudly to his grandchild when (**he**) read-3Sg a book

b. *I jajá   xerétise   tin kipéla ótan **aftí** pernúse   to ðromo.*
  The old-lady greeted-3SG the girl when **she** crossed-3SG the street.

Like adults, children of all ages accepted the same-reference reading of the null subject nearly all the time; they accepted the object reading less often; and they accepted the "other" reading even less often. In the overt subject experiment, children of all ages were like adults in accepting the object reading most of the time and the "other" reading less often, but unlike adults they over-accepted the same-reference reading of the overt pronoun until ages 10-11. Nevertheless, listening times showed that even in this condition the younger children still processed the object reading faster than the same-reference reading.
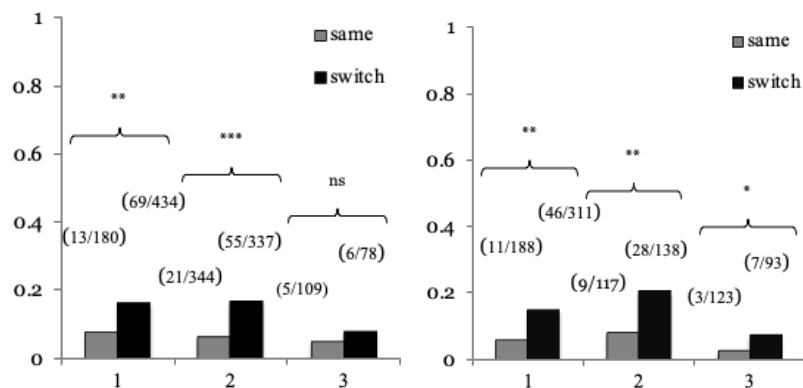
Summing up, in Italian-style pro-drop languages, children as young as 6 show sensitivity to the null/overt contrast in their own production as well as their interpretation of ambiguous pronouns. Unfortunately, not much is known about the trajectory of development before this age, after overt subjects first appear at around 2. We help fill this gap by contributing production and comprehension data for Mexican Spanish-acquiring children ages 3 to 6.

4. Distribution of null and overt subject pronouns in production

In section 2 we proposed that children acquire the contrasting antecedent preferences of null and overt subject pronouns by observing the distribution of 1st and 2nd person subjects in their input. The first question to address is whether the input even provides the necessary information—are 1st and 2nd person pronouns more likely to be overtly realized in switch-reference contexts? The second question is whether children's own production demonstrates awareness of this contrast at any time before age 6, and if so, whether they generalize it to the 3rd person.

To address the first question, we (Forsythe, Greeson & Schmitt, submitted) examined the distribution of null and overt subjects in a corpus of five mother–child dyads recorded in free-play sessions in Mexico City, Mexico. All tensed, non-imperative verbs with animate subjects were extracted and their subjects were coded for pronominal form (null, overt) and co-reference (same-, switch-reference). "Same-reference" was applied to all subject pronouns referring to the same entity as the closest preceding referential subject in the same speaker turn; "switch-reference" was applied to those referring to any other entity mentioned in the same turn[1]. The rate of overt subject realization was then compared across same- and switch-reference contexts, which we report in Figure 1.

---

[1] The decision to only examine pronoun reference within a single turn was made in order to focus on pronouns with locally available antecedents. See Forsythe, *et al.* (submitted) for details on this and other coding decisions. The fact that children perform overwhelmingly better than L2 learners on this contrast suggests that acquisition happens early and is therefore likely to be guided by the Less-is-more Principle (Newport, 1990).

Figure 1: Rate of overt 1st, 2nd, and 3rd person pronouns appearing in same- and switch-reference contexts in the speech of mothers (left) and children (right). From Forsythe *et al.*, submitted.

The first result was that 1st and 2nd person pronouns in child-directed speech were significantly more likely to be overt when they appeared in switch-reference contexts. In other words, the contrast between null and overt subjects is acquirable from 1st and 2nd person input to children. What's more, this contrast turned out to be statistically stronger in the 1st and 2nd persons, relative to the 3rd person, where the difference between same- and switch-reference contexts failed to reach significance. This shows that it is not only *possible*, but *more efficient* for children to narrow their attention to the 1st and 2nd persons, as we propose.

The second result was that children themselves also produced significantly more overt pronominal subjects in switch-reference contexts, and they did so in all three persons. That is, children have acquired this contrast and have generalized it to the 3rd person before 6 years of age.

The last important result came from an individual-level examination of pronoun distribution. Each child in this sample produced more overt subjects in switch-reference contexts than in same-reference contexts, but the contrast reached significance only for the oldest three children (4;8-5;11), as illustrated in Figure 2. This suggests that the null/overt contrast is acquired somewhere around 4 ½ years of age.
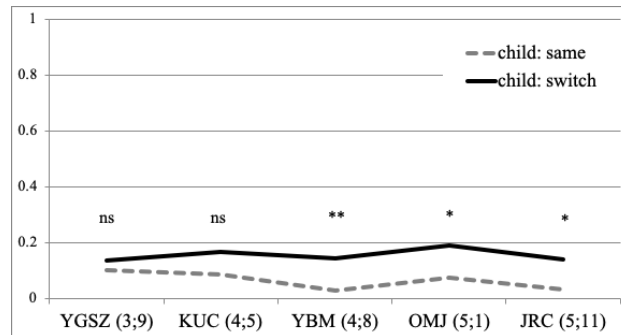
Figure 2: Rate of overt pronominal subjects produced by individual children, in same-reference and switch-reference contexts. From Forsythe *et al.*, submitted.

5. Comprehension of ambiguous null and overt subject pronouns

The third and final question to answer is whether children's sensitivity to the null/overt contrast, as shown in their production, is automatically generalized to comprehension. If, as we hypothesize, children generalize the knowledge they have acquired about the realization of 1st and 2nd person subjects to the 3rd person, then they should be able to leverage this knowledge to interpret 3rd person pronominal subjects. This predicts that children older than 4 ½ will associate null 3rd person pronouns with a same-reference reading more often than overt 3rd person pronouns. We tested this prediction using a forced-choice picture selection task probing children's preferred interpretations of grammatically ambiguous 3rd person singular subjects, as in (8). In addition to manipulating the form of the pronominal subject, we manipulated the coherence relation between the clauses containing the pronoun and its potential antecedents. This allows us to test whether sensitivity to the null/overt distinction persists across different pragmatic situations that may increase or decrease the *overall* likelihood of a same-reference reading or a switch-reference reading (Kehler 2002).

(8) a. *Juan le pega a Pedro y después ø/él se va.*
     Juan hits Pedro and then *pro*/he leaves.
   b. *Juan le pega a Pedro y por eso ø se va.*
     Juan hits Pedro and so *pro*/he leaves.

5.1. Methods
5.1.1. Subjects

A total of 40 adults (34 women) and 73 children (39 girls) ages 2;11 to 6;4 (mean: 4;6, sd: 11.5 months) completed the task. After 3 exclusions (see below), children were divided into two age groups for analysis: 40 children age 2;11-4;6 ($M = 3;8$) and 33 children age 4;7-6;4 ($M = 5;5$).

5.1.2. Design, procedure, and exclusions

Participants took a 16-item forced-choice picture selection task in a 2 (null, overt) x 2 (Version 1, Version 2) within-subjects design. In version 1 the two clauses were joined via a discourse connective expressing a temporal relation: *y después* ('and then'); in version 2 the discourse connective expressed a causal relation: *y por eso* ('and so'). Items were designed so that the causal relation would favor a switch-reference reading, while the temporal relation would still allow for a same-reference reading.

Items were created by crossing each condition with eight verb-phrase pairs (*alegrar–aplaudir*: 'cheer up–applaud,' *cantar para–bailar*: 'sing for–dance,' *perseguir–cansarse*: 'chase–get tired,' *asustar–gritar*: 'scare–yell,' *pegarle–irse*: 'hit–leave,' *pelearse con–llorar*: 'quarrel with–cry,' *tocar-reírse*: 'poke–laugh,' and *hablar con-sonreír*: 'speak to–smile'). Each participant saw every item twice, but in different conditions.

Each participant saw all four conditions, presented in blocks whose order was counterbalanced across participants. Between each block participants saw items from a separate study that used some of the same characters. Items were presented in random order within each block. Pictures were presented on a computer screen using Psychopy version 1.82.01 (Pierce 2007), and the position of the first-mentioned character (left or right side) was counterbalanced across verb-phrase pairs.

Children were read the prompts by a native speaker who recorded their responses on the computer by pressing the "4" key or "9" key, while adults listened to pre-recorded prompts and chose their responses themselves. Before beginning, participants were taught the names of the characters used in the pictures (María, Sara, Juan, Pedro) and given a 4-item recognition task. Subjects providing fewer than 2 correct answers were excluded (3 children). The remaining children gave an average of 3.37 correct responses.

5.2. Results

Figure 3 shows the rate at which participants chose the same-reference antecedent in each version of the task. For adults, 2 x 2 ANOVAs within-subjects and within-items both revealed a significant main effect of pronominal form ($F1$(1,39)=10.56, $p = 0.002$; $F2$(1,7)=12.94, $p = 0.009$) and of version ($F1$(1,39)=13.73, $p < 0.001$; $F2$(1,7)=22.23, $p = 0.002$) on the rate of same-reference choices. There was no interaction, indicating that the contrast between null and overt pronoun preferences remained constant even as the baseline preference changed from version 1 to version 2 of the experiment.
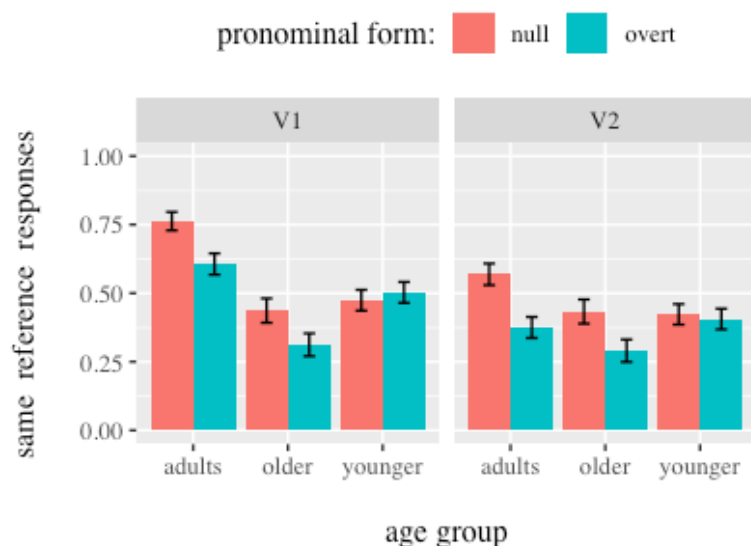
Figure 3: Proportion of subject responses by adults ($N = 40$), older children (4;7-6;4, $N = 33$), and younger children (2;11-4;6, $N = 40$). Left graph shows version 1, with temporal connective *y después* ('and then'); right graphs shows version 2, with causal connective *y por eso* ('and so').

For older children, there was a significant main effect of pronominal form ($F1(1,28)=10.0$, $p<0.005$, $F2(1,1)=178.7$, $p<0.05$) on the rate of subject antecedent choices, but no effect of version and no interaction. For younger children, there was a significant main effect of version only ($F1(1,39)=4.1$, $p<0.05$, $F2(1,1)=443.6$, $p<0.05$).

5.3. Discussion

The most basic result is that children under 6 show comprehension of the null/overt contrast—much younger than has been found using other methodologies that require metalinguistic reasoning (Shin & Cairns 2012, Sorace *et al*. 2009). The second important result is that, more precisely, children comprehend this contrast at age 4 ½, the same age at which our corpus analysis suggests they have acquired this contrast in production. These results demonstrate that children are not merely matching input frequencies; they actually show an understanding of the different antecedent preferences of null and overt subjects, a difference that they use to guide their interpretation of ambiguous pronominal subjects.

It is interesting to note that children did not share the same baseline preferences as adults. Children were overall not as inclined as adults were to choose the same-reference reading, even though younger children showed sensitivity to the contrast between temporal and causal connectives and older

children showed sensitivity to the contrast between null and overt subjects. This indicates that learning at this stage is not complete. Children must still learn to integrate additional constraints as they develop the capacity to resolve pronouns.

6. Conclusion

In this paper, we set out to identify the learning path that monolingual children follow as they acquire the constraints governing the use of null and overt subjects in pro-drop languages. Specifically, we sought to understand how children acquiring Italian-style pro-drop languages learn to probabilistically associate the null subject with reference to the preceding subject antecedent, in contrast to the overt pronominal subject. We identified the problem that this task presents to children: namely, that acquiring this contrast requires the child to identify pronoun antecedents in the first place—a non-trivial task. We proposed that children overcome this problem by first tracking the realization of 1^{st} and 2^{nd} person pronominal subjects, whose intended antecedents are much easier to identify.

Our proposal is based on the independently motivated claim by Charnavel (2017) that 1^{st}, 2^{nd} and 3^{rd} person pronouns all share the same underlying formal representation, and that their interpretive differences arise from the relative ease of deciphering the intended referent of 1^{st} and 2^{nd} person pronouns. To show that our proposal is consistent with actual child learning, we first presented results from a corpus study of ours (Forsythe *et al.*, submitted) which confirm that 1^{st} and 2^{nd} person pronominal subjects in naturalistic child-directed Spanish do indeed provide the necessary statistical signal—and in fact that the signal is stronger for these persons than for the 3^{rd} person. The same corpus study also shows that children mirror this distribution in their own production of 1^{st} and 2^{nd} person subjects, as well as 3^{rd} person pronominal subjects, and that the age of acquisition is approximately 4 ½ years.

The main thrust of this paper was to test whether this apparent sensitivity is available in comprehension, and not, for example, a product of mere frequency-matching. Using a forced-choice picture selection task, we probed children's ability to use the null/overt contrast to interpret grammatically ambiguous 3^{rd} person subject pronouns. Results indicate that children age 4 ½ to 6 (but not younger) successfully use this distinction to associate null subjects with same-reference interpretations, more often than overt subjects. This was true across environments where the clauses containing the pronoun and its two potential antecedents were linked via a temporal connective (*y después*, "and then") as well as via a causal connective (*y por eso*, "and so").

Together, these production and comprehension results make a contribution to the literature on L1 acquisition of the null/overt contrast, providing what we believe to be the first evidence that children under 6 years of age have incipient knowledge of the discourse constraints that govern the use of null and overt subjects.

References

Alonso-Ovalle, Luis; Fernández-Solera, Susana; Frazier, Lyn, & Clifton, Charles. (2002). Null vs. overt pronouns and the topic-focus articulation in Spanish. *Italian Journal of Linguistics*, *14*, 151–170.

Barbosa, Pilar (2014). pro as a Minimal n P: Towards a Unified Approach to pro-drop. *Linguistic Inquiry*, 1–78.

Bel, Aurora. (2003). The syntax of subjects in the acquisition of Spanish and Catalan. *Probus,15*(1), 1–26.

Belletti, Adriana; Bennati, Elisa, & Sorace, Antonella. (2007). Theoretical and developmental issues in the syntax of subjects: Evidence from near-native Italian. *Natural Language & Linguistic Theory*, *25*(4), 657.

Blackwell, Sarah & Quesada, Margaret Lubbers (2012). Third-Person Subjects in Native Speakers' and L2 Learners' Narratives: Testing (and Revising) the Givenness Hierarchy for Spanish. In Kimberly Geeslin and Manuel Díaz-Campos (Eds.), *Selected Proceedings of the 14th Hispanic Linguistics Symposium*. Somerville, MA: Cascadilla Proceedings Project. pp. 142–164.

Büring, Daniel (2011). Pronouns. In Claudia Maienborn & Paul Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning*. Berlin: Mouton De Gruyter. pp. 971–996.

Carminati, Maria Nella. (2002). *The processing of Italian subject pronouns* (Doctoral dissertation, ProQuest Information & Learning).

Carvalho, Ana M.; Orozco, Rafael, & Shin, Naomi L. (Eds.). (2015). *Subject pronoun expression in Spanish: A cross-dialectal perspective*. Georgetown University Press.

Charnavel, Isabelle. (2017). Presupposition failure and intended pronominal reference: Person is not so different from gender after all. *Linguistic Inquiry*, 1–35.

Enríquez, Emilia V. (1984). *El Pronombre Personal Sujeto en la Lengua Española Hablada en Madrid* (Vol. 4). Editorial CSIC-CSIC Press.

Filiaci, Francesca; Sorace, Antonella, & Carreiras, Manuel (2013). Anaphoric biases of null and overt subjects in Italian and Spanish: a cross-linguistic comparison. *Language, Cognition and Neuroscience*, 29:7, 825–843.

Forsythe, Hannah; Greeson, Daniel, & Schmitt, Cristina (submitted). How preschoolers acquire the null-overt contrast in Mexican Spanish: Evidence from production.

de la Fuente, Israel. (2015). *Putting pronoun resolution in context: The role of syntax, semantics, and pragmatics in pronoun interpretation* (Doctoral dissertation, Université Paris Diderot).

Grinstead, John. (2004). Subjects and interface delay in child Spanish and Catalan. *Language*, 40–72.

Hyams, Nina. (2012). *Language acquisition and the theory of parameters* (Vol. 3). Springer Science & Business Media.

Jegerski, Jill; VanPatten, Bill, & Keating, Gregory D. (2011). Cross-linguistic variation and the acquisition of pronominal reference in L2 Spanish. *Second Language Research*, *27*(4), 481–507.

Keating, Gregory D.; Jegerski, Jill, & Vanpatten, Bill. (2016). Online processing of subject pronouns in monolingual and heritage bilingual speakers of Mexican Spanish. *Bilingualism: Language and Cognition*, *19*(1), 36–49.

Kehler, Andrew. (2002). *Coherence, Reference, and the Theory of Grammar*. CSLI Publications.

Kratzer, Angelika. (2009). Making a pronoun: Fake indexicals as windows into the properties of pronouns. *Linguistic Inquiry*, *40*(2), 187–237.

Lastra, Yolanda, & Butragueño, Pedro Martín. (2015). Subject pronoun expression in oral Mexican Spanish. *Subject pronoun expression in Spanish: A cross-dialectal perspective*, 39–57.

Michnowicz, Jim. (2015). Subject pronoun expression in contact with Maya

in Yucatan Spanish. *Subject pronoun expression in Spanish: A cross-dialectal perspective*, 103–122.

Montrul, Silvina. (2004). Subject and object expression in Spanish heritage speakers: A case of morphosyntactic convergence. *Bilingualism: Language and cognition*, *7*(2), 125–142.

Newport, Elissa L. (1990). Maturational constraints on language learning. *Cognitive Science*, *14*(1), 11-28.

Orozco, Rafael, & Guy, Gregory. (2008). El uso variable de los pronombres sujetos: ¿Qué pasa en la costa Caribe colombiana? In *Selected Proceedings of the 4th Workshop on Spanish Sociolinguistics* (pp. 70–80). Somerville, MA: Cascadilla Proceedings Project.

Otheguy, R., Zentella, A. C., & Livert, D. (2007). Language and dialect contact in Spanish in New York: Toward the formation of a speech community. *Language*, 770–802.

Papadopoulou, Despina; Peristeri, Eleni; Plemenou, Evagelia; Marinis, Theodoros, & Tsimpli, Ianthi. (2015). Pronoun ambiguity resolution in Greek: Evidence from monolingual adults and children. *Lingua,155*, 98–120.

Peirce, J. W. (2007). PsychoPy - Psychophysics software in Python. *Journal of Neuroscience Methods*, 162 (1-2):8–13.

Sauerland, Uli. (2008). On the semantic markedness of phi-features. In *Phi Theory*, ed. Daniel Harbour, David Adger and Susana Béjar, 57-82. Oxford: Oxford University Press.

Shin, Naomi Lapidus. (2012). Variable use of Spanish subject pronouns by monolingual children in Mexico. In *Selected proceedings of the 14th Hispanic Linguistics Symposium* (pp. 130-141). Cascadilla Proceedings Project Somerville, MA.

Shin, Naomi Lapidus, & Cairns, Helen Smith. (2012). The Development of NP Selection in School-Age Children: Reference and Spanish Subject Pronouns. *Language Acquisition,19*(1), 3–38.

Shin, Naomi Lapidus, & Erker, Daniel (2015). The emergence of structured variability in morphosyntax: Childhood acquisition of Spanish subject pronouns. *Subject pronoun expression in Spanish: A cross-dialectal perspective*, 169–190.

Shin, Naomi Lapidus, & Otheguy, Ricardo. (2013). Social class and gender impacting change in bilingual settings: Spanish subject pronoun use in New York. *Language in Society*, *42*(4), 429–452.

Shin, Naomi Lapidus. (2016). Acquiring constraints on morphosyntactic variation: Children's Spanish subject pronoun expression. *Journal of child language*, *43*(4), 914–947.

Schlenker, Phillippe. (1999). *Propositional attitudes and indexicality: a cross categorial approach* (Doctoral dissertation, Massachusetts Institute of Technology).

Sorace, Antonella; Serratrice, Ludovica; Filiaci, Francesca, & Baldo, Michela (2009). Discourse conditions on subject pronoun realization: Testing the linguistic intuitions of older bilingual children. *Lingua, 119*(3), 460–477.

Valian, Virginia. (1990). Null subjects: A problem for parameter-setting models of language acquisition. *Cognition*, *35*(2), 105–122.

White, Lydia. (2011). Second language acquisition at the interfaces. *Lingua*, *121*(4), 577–590.

Yang, Charles D. (2002). *Knowledge and learning in natural language*. Oxford University Press on Demand, 2002.